



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ :

G06F 17/00

A2

(11) International Publication Number:

WO 00/54168

(43) International Publication Date: 14 September 2000 (14.09.00)

(21) International Application Number: PCT/GB00/00718

(22) International Filing Date: 1 March 2000 (01.03.00)

(30) Priority Data:

9905201.1	5 March 1999 (05.03.99)	GB
9905186.4	5 March 1999 (05.03.99)	GB
9905160.9	5 March 1999 (05.03.99)	GB
9905199.7	5 March 1999 (05.03.99)	GB
9905187.2	5 March 1999 (05.03.99)	GB

(71) Applicant (for all designated States except US): CANON KABUSHIKI KAISHA [JP/JP]; 30-2 3-chome Shimomaruko, Ohta-Ku, Tokyo (JP).

(72) Inventors; and

(75) Inventors/Applicants (for US only): CHARLESWORTH, Jason, Peter, Andrew [GB/GB]; Canon Research Centre Europe Limited, 1 Occam Court, Occam Road, Surrey Research Park, Guildford, Surrey GU2 5YJ (GB). RAJAN, Jebu, Jacob [IE/GB]; Canon Research Centre Europe Limited, 1 Occam Court, Occam Road, Surrey Research Park, Guildford, Surrey GU2 5YJ (GB). GARNER, Philip, Neil [GB/GB]; Canon Research Centre Europe Limited, 1 Occam Court, Occam Road, Surrey Research Park, Guildford, Surrey GU2 5YJ (GB).


(74) Agents: BERESFORD, Keith, Denis, Lewis et al.; Beresford & Co., 2-5 Warwick Court, High Holborn, London WC1R 5DJ (GB).

(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published

Without international search report and to be republished upon receipt of that report.

(54) Title: DATABASE ANNOTATION AND RETRIEVAL



(57) Abstract

A data structure is provided for annotating data files within a database. The annotation data comprises a phoneme and word lattice which allows the quick and efficient searching of data files within the database in response to a user's input query. The structure of the annotation data is such that it allows the input query to be made by voice and can be used for annotating various kinds of data files, such as audio data files, video data files, multimedia data files etc. The annotation data may be generated from the data files themselves or may be input by the user either from a voiced input or from a typed input.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Bélarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

DATABASE ANNOTATION AND RETRIEVAL

The present invention relates to the annotation of data files which are to be stored in a database for facilitating their subsequent retrieval. The present invention is also concerned with a system for generating the annotation data which is added to the data file and to a system for searching the annotation data in the database to retrieve a desired data file in response to a user's input query.

Databases of information are well known and suffer from the problem of how to locate and retrieve the desired information from the database quickly and efficiently. Existing database search tools allow the user to search the database using typed keywords. Whilst this is quick and efficient, this type of searching is not suitable for various kinds of databases, such as video or audio databases.

According to one aspect, the present invention aims to provide a data structure which will allow the annotation of data files within a database which will allow a quick and efficient search to be carried out in response to a user's input query.

According to one aspect, the present invention provides data defining a phoneme and word lattice for use as an annotation data for annotating data files to be stored within a database. Preferably, the data defines a plurality of nodes within the lattice and a plurality of links connecting the nodes within the lattice and further data associates a plurality of phonemes with a respective plurality of links and further data associates at least one word with at least one of said links.

- According to another aspect, the present invention provides a method of searching a database comprising the annotation data discussed above, in response to an input query by a user. The method preferably comprises the
- 5 steps of generating phoneme data and word data corresponding to the user's input query; searching the database using the word data corresponding to the user's query; selecting a portion of the data defining the phoneme and word lattice in the database for further
- 10 searching in response to the results of the word search; searching said selected portion of the database using said phoneme data corresponding to the user's input query; and outputting the search results.
- 15 According to this aspect, the present invention also provides an apparatus for searching a database which employs the annotation data discussed above for annotating data files therein. The apparatus preferably comprises means for generating phoneme data and word data
- 20 corresponding to a user's input query; means for searching the database using the word data corresponding to the user's input query to identify similar words within the database; means for selecting a portion of the annotation data in the database for further searching in
- 25 response to the results of the word search; means for searching the selected portion using the phoneme data corresponding to the user's input query; and means for outputting the search results.
- 30 The phoneme and annotation data for a data file may be generated from the data file itself or from a typed or spoken annotation input by the user.

Exemplary embodiments of the present invention will now

35 be described with reference to the accompanying figures,

in which:

Figure 1 is a schematic view of a computer which is programmed to operate an embodiment of the present invention;

Figure 2 is a block diagram showing a phoneme and word annotator unit which is operable to generate phoneme and word annotation data for appendage to a data file;

Figure 3 is a block diagram illustrating one way in which the phoneme and word annotator can generate the annotation data from an input video data file;

Figure 4a is a schematic diagram of a phoneme lattice for an example audio string from the input video data file;

Figure 4b is a schematic diagram of a word and phoneme lattice embodying one aspect of the present invention, for an example audio string from the input video data file;

Figure 5 is a schematic block diagram of a user's terminal which allows the user to retrieve information from the database by a voice query;

Figure 6a is a flow diagram illustrating part of the flow control of the user terminal shown in Figure 5;

Figure 6b is a flow diagram illustrating the remaining part of the flow control of the user terminal shown in Figure 5;

Figure 7 is a flow diagram illustrating the way in which a search engine forming part of the user's terminal

carries out a phoneme search within the database;

Figure 8 is a schematic diagram illustrating the form of
a phoneme string and four M-GRAMS generated from the
5 phoneme string;

Figure 9 is a plot showing two vectors and the angle
between the two vectors;

10 Figure 10 is a schematic diagram of a pair of word and
phoneme lattices, for example audio strings from two
speakers;

Figure 11 is a schematic block diagram illustrating a
15 user terminal which allows the annotation of a data file
with annotation data generated from an audio signal input
from a user;

Figure 12 is a schematic diagram of phoneme and word
20 lattice annotation data which is generated for an example
utterance input by the user for annotating a data file;

Figure 13 is a schematic block diagram illustrating a
user terminal which allows the annotation of a data file
25 with annotation data generated from a typed input from a
user;

Figure 14 is a schematic diagram of phoneme and word
lattice annotation data which is generated for a typed
30 input by the user for annotating a data file;

Figure 15 is a block schematic diagram showing the form
of a document annotation system;

35 Figure 16 is a block schematic diagram of an alternative

document annotation system;

Figure 17 is a block schematic diagram of another document annotation system;

5

Figure 18 is a schematic block diagram illustrating a user terminal which is operable to access a database located on a remote server via a data network in response to an input utterance by the user;

10

Figure 19 is a schematic block diagram of a user terminal which allows a user to access a database located in a remote server in response to an input utterance from the user;

15

Figure 20 is a schematic block diagram of a user terminal which allows a user to access a database by a typed input query; and

20

Figure 21 is a schematic block diagram illustrating the way in which a phoneme and word lattice can be generated from script data contained within a video data file.

25

Embodiments of the present invention can be implemented using dedicated hardware circuits, but the embodiment to be described is implemented in computer software or code, which is run in conjunction with processing hardware such as a personal computer, work station, photocopier, facsimile machine, personal digital assistant (PDA) or the like.

30

Figure 1 shows a personal computer (PC) 1 which is programmed to operate an embodiment of the present invention. A keyboard 3, a pointing device 5, a microphone 7 and a telephone line 9 are connected to the

35

PC 1 via an interface 11. The keyboard 3 and pointing device 5 enable the system to be controlled by a user. The microphone 7 converts acoustic speech signals from the user into equivalent electrical signals and supplies them to the PC 1 for processing. An internal modem and speech receiving circuit (not shown) is connected to the telephone line 9 so that the PC 1 can communicate with, for example, a remote computer or with a remote user.

10. The programme instructions which make the PC 1 operate in accordance with the present invention may be supplied for use with an existing PC 1 on, for example, a storage device such as a magnetic disc 13, or by downloading the software from the Internet (not shown) via the internal modem and telephone line 9.

DATA FILE ANNOTATION

Figure 2 is a block diagram illustrating the way in which annotation data 21 for an input data file 23 is generated in this embodiment by a phoneme and word annotating unit 25. As shown, the generated phoneme and word annotation data 21 is then combined with the data file 23 in the data combination unit 27 and the combined data file output thereby is input to the database 29. In this embodiment, the annotation data 21 comprises a combined phoneme (or phoneme like) and word lattice which allows the user to retrieve information from the database by a voice query. As those skilled in the art will appreciate, the data file 23 can be any kind of data file, such as, a video file, an audio file, a multimedia file etc.

A system has been proposed to generate N-Best word lists for an audio stream as annotation data by passing the audio data from a video data file through an automatic

speech recognition unit. However, such word-based systems suffer from a number of problems. These include (i) that state of the art speech recognition systems still make basic mistakes in recognition; (ii) that state of the art automatic speech recognition systems use a dictionary of perhaps 20,000 to 100,000 words and cannot produce words outside that vocabulary; and (iii) that the production of N-Best lists grows exponentially with the number of hypothesis at each stage, therefore resulting in the annotation data becoming prohibitively large for long utterances.

The first of these problems may not be that significant if the same automatic speech recognition system is used to generate the annotation data and to subsequently retrieve the corresponding data file, since the same decoding error could occur. However, with advances in automatic speech recognition systems being made each year, it is likely that in the future the same type of error may not occur, resulting in the inability to be able to retrieve the corresponding data file at that later date. With regard to the second problem, this is particularly significant in video data applications, since users are likely to use names and places (which may not be in the speech recognition dictionary) as input query terms. In place of these names, the automatic speech recognition system will typically replace the out of vocabulary words with a phonetically similar word or words within the vocabulary, often corrupting nearby decodings. This can also result in the failure to retrieve the required data file upon subsequent request.

In contrast, with the proposed phoneme and word lattice annotation data, a quick and efficient search using the word data in the database 29 can be carried out and, if

this fails to provide the required data file, then a further search using the more robust phoneme data can be performed. The phoneme and word lattice is an acyclic directed graph with a single entry point and a single exit point. It represents different parses of the audio stream within the data file. It is not simply a sequence of words with alternatives since each word does not have to be replaced by a single alternative, one word can be substituted for two or more words or phonemes, and the whole structure can form a substitution for one or more words or phonemes. Therefore, the density of data within the phoneme and word lattice essentially remains linear throughout the audio data, rather than growing exponentially as in the case of the N-Best technique discussed above. As those skilled in the art of speech recognition will realise, the use of phoneme data is more robust, because phonemes are dictionary independent and allow the system to cope with out of vocabulary words, such as names, places, foreign words etc. The use of phoneme data is also capable of making the system future proof, since it allows data files which are placed into the database to be retrieved even when the words were not understood by the original automatic speech recognition system.

The way in which this phoneme and word lattice annotation data can be generated for a video data file will now be described with reference to Figure 3. As shown, the video data file 31 comprises video data 31-1, which defines the sequence of images forming the video sequence and audio data 31-2, which defines the audio which is associated with the video sequence. As is well known, the audio data 31-2 is time synchronised with the video data 31-1 so that, in use, both the video and audio data are supplied to the user at the same time.

As shown in Figure 3, in this embodiment, the audio data 31-2 is input to an automatic speech recognition unit 33, which is operable to generate a phoneme lattice corresponding to the stream of audio data 31-2. Such an
5 automatic speech recognition unit 33 is commonly available in the art and will not be described in further detail. The reader is referred to, for example, the book entitled 'Fundamentals of Speech Recognition' by Lawrence Rabiner and Bing-Hwang Juang and, in particular, to
10 pages 42 to 50 thereof, for further information on this type of speech recognition system.

Figure 4a illustrates the form of the phoneme lattice data output by the speech recognition unit 33, for the
15 input audio corresponding to the phrase '...tell me about Jason...'. As shown, the automatic speech recognition unit 33 identifies a number of different possible phoneme strings which correspond to this input audio utterance. For example, the speech recognition system considers that
20 the first phoneme in the audio string is either a /t/ or a /d/. As is well known in the art of speech recognition, these different possibilities can have their own weighting which is generated by the speech recognition unit 33 and is indicative of the confidence
25 of the speech recognition unit's output. For example, the phoneme /t/ may be given a weighting of 0.9 and the phoneme /d/ may be given a weighting of 0.1, indicating that the speech recognition system is fairly confident that the corresponding portion of audio represents the
30 phoneme /t/, but that it still may be the phoneme /d/. In this embodiment, however, this weighting of the phonemes is not performed.

As shown in Figure 3, the phoneme lattice data 35 output
35 by the automatic speech recognition unit 33 is input to

a word decoder 37 which is operable to identify possible words within the phoneme lattice data 35. In this embodiment, the words identified by the word decoder 37 are incorporated into the phoneme lattice data structure.

5 For example, for the phoneme lattice shown in Figure 4a, the word decoder 37 identifies the words 'tell', 'dell', 'term', 'me', 'a', 'boat', 'about', 'chase' and 'sun'.

As shown in Figure 4b, these identified words are added to the phoneme lattice data structure output by the
10 speech recognition unit 33, to generate a phoneme and word lattice data structure which forms the annotation data 31-3. This annotation data 31-3 is then combined with the video data file 31 to generate an augmented video data file 31' which is then stored in the database

15 29. As those skilled in the art will appreciate, in a similar way to the way in which the audio data 31-2 is time synchronised with the video data 31-1, the annotation data 31-3 is also time synchronised and associated with the corresponding video data 31-1 and
20 audio data 31-2, so that a desired portion of the video and audio data can be retrieved by searching for and locating the corresponding portion of the annotation data 31-3.

25 In this embodiment, the annotation data 31-3 stored in the database 29 has the following general form:

HEADER

- time of start
- flag if word if phoneme if mixed
- 30 - time index associating the location of blocks of annotation data within memory to a given time point.
- word set used (i.e. the dictionary)
- phoneme set used
- 35 - the language to which the vocabulary

11

pertains

Block(i) $i = 0, 1, 2, \dots$ node N_j $j = 0, 1, 2, \dots$

- time offset of node from start of block

5 - phoneme links (k) $k = 0, 1, 2, \dots$ offset to node $N_j = N_k - N_j$ (N_k is node to which link k extends) or if N_k is in10 block (i+1) offset to node $N_j = N_k + N_b - N_j$ (where N_b is the number of nodes in block (i))

phoneme associated with link (k)

- word links (l) $l = 0, 1, 2, \dots$ offset to node $N_j = N_l - N_j$ (N_l is node to which link l extends) or if N_l is in15 block (i+1) offset to node $N_j = N_l + N_b - N_j$ (where N_b is the number of nodes in block (i))

word associated with link (l)

20 The time of start data in the header can identify the time and date of transmission of the data. For example, if the video file is a news broadcast, then the time of start may include the exact time of the broadcast and the date on which it was broadcast.

25

The flag identifying if the annotation data is word annotation data, phoneme annotation data or if it is mixed is provided since not all the data files within the database will include the combined phoneme and word

30 lattice annotation data discussed above, and in this case, a different search strategy would be used to search this annotation data.

In this embodiment, the annotation data is divided into

35 blocks in order to allow the search to jump into the

middle of the annotation data for a given audio data stream. The header therefore includes a time index which associates the location of the blocks of annotation data within the memory to a given time offset between the time of start and the time corresponding to the beginning of the block.

The header also includes data defining the word set used (i.e. the dictionary), the phoneme set used and the language to which the vocabulary pertains. The header may also include details of the automatic speech recognition system used to generate the annotation data and any appropriate settings thereof which were used during the generation of the annotation data.

15

The blocks of annotation data then follow the header and identify, for each node in the block, the time offset of the node from the start of the block, the phoneme links which connect that node to other nodes by phonemes and word links which connect that node to other nodes by words. Each phoneme link and word link identifies the phoneme or word which is associated with the link. They also identify the offset to the current node. For example, if node N_{50} is linked to node N_{55} by a phoneme link, then the offset to node N_{50} is 5. As those skilled in the art will appreciate, using an offset indication like this allows the division of the continuous annotation data into separate blocks.

In an embodiment where an automatic speech recognition unit outputs weightings indicative of the confidence of the speech recognition units output, these weightings or confidence scores would also be included within the data structure. In particular, a confidence score would be provided for each node which is indicative of the

35

confidence of arriving at the node and each of the phoneme and word links would include a transition score depending upon the weighting given to the corresponding phoneme or word. These weightings would then be used to control the search and retrieval of the data files by discarding those matches which have a low confidence score.

DATA FILE RETRIEVAL

Figure 5 is a block diagram illustrating the form of a user terminal 59 which can be used to retrieve the annotated data files from the database 29. This user terminal 59 may be, for example, a personal computer, hand held device or the like. As shown, in this embodiment, the user terminal 59 comprises the database 29 of annotated data files, an automatic speech recognition unit 51, a search engine 53, a control unit 55 and a display 57. In operation, the automatic speech recognition unit 51 is operable to process an input voice query from the user 39 received via the microphone 7 and the input line 61 and to generate therefrom corresponding phoneme and word data. This data may also take the form of a phoneme and word lattice, but this is not essential. This phoneme and word data is then input to the control unit 55 which is operable to initiate an appropriate search of the database 29 using the search engine 53. The results of the search, generated by the search engine 53, are then transmitted back to the control unit 55 which analyses the search results and generates and displays appropriate display data to the user via the display 57.

Figures 6a and 6b are flow diagrams which illustrate the way in which the user terminal 59 operates in this embodiment. In step s1, the user terminal 59 is in an

idle state and awaits an input query from the user 39. Upon receipt of an input query, the phoneme and word data for the input query is generated in step s3 by the automatic speech recognition unit 51. The control unit
5 55 then instructs the search engine 53, in step s5, to perform a search in the database 29 using the word data generated for the input query. The word search employed in this embodiment is the same as is currently being used in the art for typed keyword searches, and will not be
10 described in more detail here. If in step s7, the control unit 55 identifies from the search results, that a match for the user's input query has been found, then it outputs the search results to the user via the display 57.

15 In this embodiment, the user terminal 59 then allows the user to consider the search results and awaits the user's confirmation as to whether or not the results correspond to the information the user requires. If they are, then
20 the processing proceeds from step s11 to the end of the processing and the user terminal 59 returns to its idle state and awaits the next input query. If, however, the user indicates (by, for example, inputting an appropriate voice command) that the search results do not correspond
25 to the desired information, then the processing proceeds from step s11 to step s13, where the search engine 53 performs a phoneme search of the database 29. However, in this embodiment, the phoneme search performed in step s13 is not of the whole database 29, since this could
30 take several hours depending on the size of the database 29.

Instead, the phoneme search performed in step s13 uses the results of the word search performed in step s5 to
35 identify one or more portions within the database which

may correspond to the user's input query. The way in which the phoneme search performed in step s13 is performed in this embodiment, will be described in more detail later. After the phoneme search has been performed, the control unit 55 identifies, in step s15, if a match has been found. If a match has been found, then the processing proceeds to step s17 where the control unit 55 causes the search results to be displayed to the user on the display 57. Again, the system then awaits the user's confirmation as to whether or not the search results correspond to the desired information. If the results are correct, then the processing passes from step s19 to the end and the user terminal 59 returns to its idle state and awaits the next input query. If however, the user indicates that the search results do not correspond to the desired information, then the processing proceeds from step s19 to step s21, where the control unit 55 is operable to ask the user, via the display 57, whether or not a phoneme search should be performed of the whole database 29. If in response to this query, the user indicates that such a search should be performed, then the processing proceeds to step s23 where the search engine performs a phoneme search of the entire database 29.

On completion of this search, the control unit 55 identifies, in step s25, whether or not a match for the user's input query has been found. If a match is found, then the processing proceeds to step s27 where the control unit 55 causes the search results to be displayed to the user on the display 57. If the search results are correct, then the processing proceeds from step s29 to the end of the processing and the user terminal 59 returns to its idle state and awaits the next input query. If, on the other hand, the user indicates that the

search results still do not correspond to the desired information, then the processing passes to step s31 where the control unit 55 queries the user, via the display 57, whether or not the user wishes to redefine or amend the search query. If the user does wish to redefine or amend the search query, then the processing returns to step s3 where the user's subsequent input query is processed in a similar manner. If the search is not to be redefined or amended, then the search results and the user's initial input query are discarded and the user terminal 59 returns to its idle state and awaits the next input query.

PHONEME SEARCH

As mentioned above, in steps s13 and s23, the search engine 53 compares the phoneme data of the input query with the phoneme data in the phoneme and word lattice annotation data stored in the database 29. Various techniques can be used including standard pattern matching techniques such as dynamic programming, to carry out this comparison. In this embodiment, a technique which we refer to as M-GRAMS is used. This technique was proposed by Ng, K. and Zue, V.W. and is discussed in, for example, the paper entitled "Subword unit representations for spoken document retrieval" published in the proceedings of Eurospeech 1997.

The problem with searching for individual phonemes is that there will be many occurrences of each phoneme within the database. Therefore, an individual phoneme on its own does not provide enough discriminability to be able to match the phoneme string of the input query with the phoneme strings within the database. Syllable sized units, however, are likely to provide more discriminability, although they are not easy to identify.

The M-GRAM technique presents a suitable compromise between these two possibilities and takes overlapping fixed size fragments, or M-GRAMS, of the phoneme string to provide a set of features. This is illustrated in Figure 8, which shows part of an input phoneme string having phonemes a, b, c, d, e, and f, which are split into four M-GRAMS (a, b, c), (b, c, d), (c, d, e) and (d, e, f). In this illustration, each of the four M-GRAMS comprises a sequence of three phonemes which is unique and represents a unique feature (f_i) which can be found within the input phoneme string.

Therefore, referring to Figure 7, the first step s51 in performing the phoneme search in step s13 shown in Figure 6, is to identify all the different M-GRAMS which are in the input phoneme data and their frequency of occurrence. Then, in step s53, the search engine 53 determines the frequency of occurrence of the identified M-GRAMS in the selected portion of the database (identified from the word search performed in step s5 in Figure 6). To illustrate this, for a given portion of the database and for the example M-GRAMS illustrated in Figure 8, this yields the following table of information:

M-GRAM (feature (f_i))	Input phoneme string frequency of occurrence (q)	Phoneme string of selected portion of database (a)
M_1	1	0
M_2	2	2
M_3	3	2
M_4	1	1

Next, in step s55, the search engine 53 calculates a similarity score representing a similarity between the phoneme string of the input query and the phoneme string of the selected portion from the database. In this embodiment, this similarity score is determined using a cosine measure using the frequencies of occurrence of the identified M-GRAMS in the input query and in the selected portion of the database as vectors. The philosophy behind this technique is that if the input phoneme string is similar to the selected portion of the database phoneme string, then the frequency of occurrence of the M-GRAM features will be similar for the two phoneme strings. Therefore, if the frequencies of occurrence of the M-GRAMS are considered to be vectors (i.e. considering the second and third columns in the above table as vectors), then if there is a similarity between the input phoneme string and the selected portion of the database, then the angle between these vectors should be small. This is illustrated in Figure 9 for two-dimensional vectors \underline{a} and \underline{q} , with the angle between the vectors given as θ . In the example shown in Figure 8, the vectors \underline{a} and \underline{q} will be four dimensional vectors and the similarity score can be calculated from:

$$SCORE = \cos \theta = \frac{\underline{a} \cdot \underline{q}}{|\underline{a}| |\underline{q}|} \quad (1)$$

This score is then associated with the current selected portion of the database and stored until the end of the search. In some applications, the vectors used in the calculation of the cosine measure will be the logarithm of these frequencies of occurrences, rather than the frequencies of occurrences themselves.

The processing then proceeds to step s57 where the search engine 53 identifies whether or not there are any more selected portions of phoneme strings from the database 29. If there are, then the processing returns to step s53 where a similar procedure is followed to identify the score for this portion of the database. If there are no more selected portions, then the searching ends and the processing returns to step s15 shown in Figure 6, where the control unit considers the scores generated by the search engine 53 and identifies whether or not there is a match by, for example, comparing the calculated scores with a predetermined threshold value.

As those skilled in the art will appreciate, a similar matching operation will be performed in step s23 shown in Figure 6. However, since the entire database is being searched, this search is carried out by searching each of the blocks discussed above in turn.

20 ALTERNATIVE EMBODIMENTS

As those skilled in the art will appreciate, this type of phonetic and word annotation of data files in a database provides a convenient and powerful way to allow a user to search the database by voice. In the illustrated embodiment, a single audio data stream was annotated and stored in the database for subsequent retrieval by the user. As those skilled in the art will appreciate, when the input data file corresponds to a video data file, the audio data within the data file will usually include audio data for different speakers. Instead of generating a single stream of annotation data for the audio data, separate phoneme and word lattice annotation data can be generated for the audio data of each speaker. This may be achieved by identifying, from the pitch or from another distinguishing feature of the speech signals, the

audio data which corresponds to each of the speakers and then by annotating the different speaker's audio separately. This may also be achieved if the audio data was recorded in stereo or if an array of microphones were used in generating the audio data, since it is then possible to process the audio data to extract the data for each speaker.

Figure 10 illustrates the form of the annotation data in such an embodiment, where a first speaker utters the words "... this so" and the second speaker replies "yes". As illustrated, the annotation data for the different speakers' audio data are time synchronised, relative to each other, so that the annotation data is still time synchronised to the video and audio data within the data file. In such an embodiment, the header information in the data structure should preferably include a list of the different speakers within the annotation data and, for each speaker, data defining that speaker's language, accent, dialect and phonetic set, and each block should identify those speakers that are active in the block.

In the above embodiments, a speech recognition system was used to generate the annotation data for annotating a data file in the database. As those skilled in the art will appreciate, other techniques can be used to generate this annotation data. For example, a human operator can listen to the audio data and generate a phonetic and word transcription to thereby manually generate the annotation data.

In the above embodiments, the annotation data was generated from audio stored in the data file itself. As those skilled in the art will appreciate, other techniques can be used to input the annotation data.

Figure 11 illustrates the form of a user terminal 59 which allows a user to input voice annotation data via the microphone 7 for annotating a data file 91 which is to be stored in the database 29. In this embodiment, the data file 91 comprises a two dimensional image generated by, for example, a camera. The user terminal 59 allows the user 39 to annotate the 2D image with an appropriate annotation which can be used subsequently for retrieving the 2D image from the database 29. In this embodiment, the input voice annotation signal is converted, by the automatic speech recognition unit 51, into phoneme and word lattice annotation data which is passed to the control unit 55. In response to the user's input, the control unit 55 retrieves the appropriate 2D file from the database 29 and appends the phoneme and word annotation data to the data file 91. The augmented data file is then returned to the database 29. During this annotating step, the control unit 55 is operable to display the 2D image on the display 57 so that the user can ensure that the annotation data is associated with the correct data file 91.

The automatic speech recognition unit 51 generates the phoneme and word lattice annotation data by (i) generating a phoneme lattice for the input utterance; (ii) then identifying words within the phoneme lattice; and (iii) finally by combining the two. Figure 12 illustrates the form of the phoneme and word lattice annotation data generated for the input utterance "picture of the Taj-Mahal". As shown, the automatic speech recognition unit identifies a number of different possible phoneme strings which correspond to this input utterance. As shown in Figure 12, the words which the automatic speech recognition unit 51 identifies within the phoneme lattice are incorporated into the phoneme

lattice data structure. As shown, for the example phrase, the automatic speech recognition unit 51 identifies the words "picture", "of", "off", "the", "other", "ta", "tar", "jam", "ah", "hal", "ha" and "al". The control unit 55 is then operable to add this annotation data to the 2D image data file 91 which is then stored in a database 29.

As those skilled in the art will appreciate, this embodiment can be used to annotate any kind of image such as x-rays of patients, 3D videos of, for example, NMR scans, ultrasound scans etc. It can also be used to annotate one-dimensional data, such as audio data or seismic data.

In the above embodiment, a data file was annotated from a voiced annotation. As those skilled in the art will appreciate, other techniques can be used to input the annotation. For example, Figure 13 illustrates the form of a user terminal 59 which allows a user to input typed annotation data via the keyboard 3 for annotating a data file 91 which is to be stored in a database 29. In this embodiment, the typed input is converted, by the phonetic transcription unit 75, into the phoneme and word lattice annotation data (using an internal phonetic dictionary (not shown)) which is passed to the control unit 55. In response to the user's input, the control unit 55 retrieves the appropriate 2D file from the database 29 and appends the phoneme and word annotation data to the data file 91. The augmented data file is then returned to the database 29. During this annotating step, the control unit 55 is operable to display the 2D image on the display 57 so that the user can ensure that the annotation data is associated with the correct data file 91.

Figure 14 illustrates the form of the phoneme and word lattice annotation data generated for the input utterance "picture of the Taj-Mahal". As shown in Figure 2, the phoneme and word lattice is an acyclic directed graph with a single entry point and a single exit point. It represents different parses of the user's input. As shown, the phonetic transcription unit 75 identifies a number of different possible phoneme strings which correspond to the typed input.

Figure 15 is a block diagram illustrating a document annotation system. In particular, as shown in Figure 15, a text document 101 is converted into an image data file by a document scanner 103. The image data file is then passed to an optical character recognition (OCR) unit 105 which converts the image data of the document 101 into electronic text. This electronic text is then supplied to a phonetic transcription unit 107 which is operable to generate phoneme and word annotation data 109 which is then appended to the image data output by the scanner 103 to form a data file 111. As shown, the data file 111 is then stored in the database 29 for subsequent retrieval. In this embodiment, the annotation data 109 comprises the combined phoneme and word lattice described above which allows the user to subsequently retrieve the data file 111 from the database 29 by a voice query.

Figure 16 illustrates a modification to the document annotation system shown in Figure 15. The difference between the system shown in Figure 16 and the system shown in Figure 15 is that the output of the optical character recognition unit 105 is used to generate the data file 113, rather than the image data output by the scanner 103. The rest of the system shown in Figure 16 is the same as that shown in Figure 15 and will not be

described further.

Figure 17 shows a further modification to the document annotation system shown in Figure 15. In the embodiment shown in Figure 17, the input document is received by a facsimile unit 115 rather than a scanner 103. The image data output by the facsimile unit is then processed in the same manner as the image data output by the scanner 103 shown in Figure 15, and will not be described again.

In the above embodiment, a phonetic transcription unit 107 was used for generating the annotation data for annotating the image or text data. As those skilled in the art will appreciate, other techniques can be used. For example, a human operator can manually generate this annotation data from the image of the document itself.

In the above embodiment, the database 29 and the automatic speech recognition unit were both located within the user terminal 59. As those skilled in the art will appreciate, this is not essential. Figure 18 illustrates an embodiment in which the database 29 and the search engine 53 are located in a remote server 60 and in which the user terminal 59 accesses and controls data files in the database 29 via the network interface units 67 and 69 and a data network 68 (such as the Internet). In operation, the user inputs a voice query via the microphone 7 which is converted into phoneme and word data by the automatic speech recognition unit 51. This data is then passed to the control unit which controls the transmission of this phoneme and word data over the data network 68 to the search engine 53 located within the remote server 60. The search engine 53 then carries out the search in accordance with the received phoneme and word data or controls the manipulation of the

data files (for example to control the playing, forwarding or rewinding of a video file) in accordance with the received phoneme and word data. The data retrieved from the database 29 or other data relating to the search is then transmitted back, via the data network 68, to the control unit 55 which controls the display of appropriate data on the display 57 for viewing by the user 39. In this way it is possible to retrieve and control data files in the remote server 60 without using significant computer resources in the server (since it is the user terminal 59 which converts the input speech into the phoneme and word data).

In addition to locating the database 29 and the search engine 53 in the remote server 60, it is also possible to locate the automatic speech recognition unit 51 in the remote server 60. Such an embodiment is shown in Figure 19. As shown in this embodiment, the input voice query from the user is passed via input line 61 to a speech encoding unit 73 which is operable to encode the speech for efficient transfer through the data network 68. The encoded data is then passed to the control unit 55 which transmits the data over the network 68 to the remote server 60, where it is processed by the automatic speech recognition unit 51. The phoneme and word data generated by the speech recognition unit 51 for the input query is then passed to the search engine 53 for use in searching and controlling data files in the database 29. Appropriate data retrieved by the search engine 53 is then passed, via the network interface 69 and the network 68, back to the user terminal 59. This data received back from the remote server is passed via the network interface unit 67 to the control unit 55 which generates and displays appropriate data on the display 57 for viewing by the user.

In the above embodiments, the user inputs his query by voice. Figure 20 shows an alternative embodiment in which the user inputs the query via the keyboard 3. As shown, the text input via the keyboard 3 is passed to
5 phonetic transcription unit 75 which is operable to generate a corresponding phoneme string from the input text. This phoneme string together with the words input via the keyboard 3 are then passed to the control unit 55 which initiates a search of database using the search
10 engine 53. The way in which this search is carried out is the same as in the first embodiment and will not, therefore, be described again. As with the other embodiments discussed above, the phonetic transcription unit 75, search engine 53 and/or the database 29 may all
15 be located in a remote server.

In the first embodiment, the audio data from the data file 31 was passed through an automatic speech recognition unit in order to generate the phoneme
20 annotation data. In some situations, a transcript of the audio data will be present in the data file. Such an embodiment is illustrated in Figure 21. In this embodiment, the data file 81 represents a digital video file having video data 81-1, audio data 81-2 and script
25 data 81-3 which defines the lines for the various actors in the video film. As shown, the script data 81-3 is passed through a text to phoneme converter 83, which generates phoneme lattice data 85 using a stored dictionary which translates words into possible sequences
30 of phonemes. This phoneme lattice data 85 is then combined with the script data 81-3 to generate the above described phoneme and word lattice annotation data 81-4. This annotation data is then added to the data file 81 to generate an augmented data file 81' which is then added
35 to the database 29. As those skilled in the art will

appreciate, this embodiment facilitates the generation of separate phoneme and word lattice annotation data for the different speakers within the video data file, since the script data usually contains indications of who is talking. The synchronisation of the phoneme and word lattice annotation data with the video and audio data can then be achieved by performing a forced time alignment of the script data with the audio data using an automatic speech recognition system (not shown).

In the above embodiments, a phoneme (or phoneme-like) and word lattice was used to annotate a data file. As those skilled in the art of speech recognition and speech processing will realise, the word "phoneme" in the description and claims is not limited to its linguistic meaning but includes the various sub-word units that are identified and used in standard speech recognition systems.

CLAIMS:

1. Data defining a phoneme and word lattice, the data comprising:

5 data for defining a plurality of nodes within the lattice and a plurality of links connecting the nodes within the lattice;

data associating a plurality of phonemes with a respective plurality of links; and

10 data associating at least one word with at least one of said links.

2. Data according to any preceding claim, wherein said data defining said phoneme and word lattice is arranged
15 in blocks of nodes.

3. Data according to claim 1, further comprising data defining time stamp information for each of said nodes.

20 4. Data according to claim 3, arranged in blocks of equal time duration.

5. Data according to claim 2 or 4, further comprising data defining each blocks location within said database.

25 6. Data according to claim 3 or any claim dependent thereon, wherein said data defining a phoneme and word lattice is associated with further data defining a time sequential signal, and wherein said time stamp
30 information is time synchronised with said time sequential signal.

7. Data according to claim 6, wherein said further data defines an audio and/or video signal. .

8. Data according to claim 7, wherein said further data defines at least speech data and wherein said data defining said phoneme and word lattice is derived from said further data.

9. Data according to claim 8, wherein said speech data comprises audio data and wherein said data defining said phoneme and word lattice is derived by passing said audio signal through an automatic speech recognition system.

10. Data according to claim 8 or 9, wherein said speech data defines the parol of a plurality of speakers, and wherein said data defines a separate phoneme and word lattice for the parol of each speaker.

11. Data according to any preceding claim, further comprising data defining a weighting for the phonemes and/or words associated with said links.

12. Data according to any preceding claim, wherein at least one of said nodes is connected to a plurality of other nodes by a plurality of links.

13. Data according to claim 12, wherein at least one of said plurality of links connecting said node to said plurality of other nodes is associated with a phoneme and wherein at least one of said links connecting said node to said plurality of other nodes is associated with a word.

14. A method of searching a database comprising data according to any preceding claim, in response to an input query, the method comprising the steps of:

generating phoneme data and/or word data corresponding to the input query;

searching the phoneme and word lattice using the phoneme and/or word data generated for the input query; and

5 outputting search results in dependence upon the results of said searching step.

15. A method according to claim 14, wherein said searching step comprises the steps of:

10 (i) searching the phoneme and word lattice using the word data generated for the user's input query to identify similar words within the phoneme and word lattice;

15 (ii) selecting one or more portions of the phoneme and word lattice for further searching in response to the results of said word search; and

(iii) searching said one or more selected portions of the phoneme and word lattice using the phoneme data generated for the user's input query.

20 16. A method according to claim 15, wherein the results of the word search are output to the user before the phoneme search is performed on the selected portions of the database.

25 17. A method according to claim 16, wherein said phoneme search is only performed in response to a further input by the user in response to the outputting of the results from the word search.

30 18. A method according to any of claims 15 to 17, wherein said phoneme search is carried out by identifying a number of features within the phoneme sequence corresponding to the user's input query and identifying similar features within the data defining said phoneme
35 lattice within the database.

19. A method according to claim 18, wherein each of said features represents a unique sequence of phonemes within the phoneme data of the user's input query.

5 20. A method according to claim 19, wherein said phoneme search employs a cosine measure to indicate the similarity between the phoneme data corresponding to the user's input query and the phoneme data within the database.

10 21. A method according to any of claims 14 to 20, wherein said search results are output to a display.

15 22. A method according to any of claims 14 to 21, wherein said input query by the user is input by voice, and wherein said step of generating phoneme data and word data employs an automatic speech recognition system.

20 23. A method according to any of claims 14 to 21, wherein said input query is a typed input and wherein said step of generating phoneme data and word data employs a text-to-phoneme converter.

25 24. An apparatus for searching a database comprising data according to any of claims 1 to 13, in response to an input query, the apparatus comprising:

means for generating phoneme data and/or word data corresponding to the input query;

30 means for searching the phoneme and word lattice using the phoneme and/or word data generated for the input query; and

means for outputting search results in dependence upon the output from said searching means.

35 25. An apparatus according to claim 24, wherein said

searching means comprises:

(i) means for searching the phoneme and word lattice using the word data generated for the user's input query to identify similar words within the phoneme and word lattice;

(ii) means for selecting one or more portions of the phoneme and word lattice for further searching in response to the results of said word search; and

(iii) means for searching said one or more selected portions of the phoneme and word lattice using the phoneme data generated for the user's input query.

26. An apparatus according to claim 25, wherein said output means is operable to output the results of the word search to the user before the phoneme search is performed on the selected portions of the database.

27. An apparatus according to claim 26, wherein said phoneme search is only performed in response to a further input by the user in response to the outputting of the results from the word search.

28. An apparatus according to any of claims 25 to 27, wherein said phoneme search is carried out by identifying a number of features within the phoneme sequence corresponding to the user's input query and identifying similar features within the data defining said phoneme lattice within the database.

29. An apparatus according to claim 28, wherein each of said features represents a unique sequence of phonemes within the phoneme data of the user's input query.

30. An apparatus according to claim 29, wherein said phoneme search employs a cosine measure to indicate the

similarity between the phoneme data corresponding to the user's input query and the phoneme data within the database.

5 31. An apparatus according to any of claims 24 to 30, wherein said output means comprises a display.

10 32. An apparatus according to any of claims 24 to 31, wherein said input query by the user is a voice query, and wherein said means for generating phoneme data and word data comprises an automatic speech recognition system which is operable to generate said phoneme data and a word decoder which is operable to generate said word data.

15 33. An apparatus according to any of claims 24 to 31, wherein said input query is a typed query and wherein said means for generating phoneme data and word data comprises a text-to-phoneme converter which is operable
20 to generate said phoneme data.

34. An apparatus for generating annotation data for use in annotating a data file comprising audio data, the apparatus comprising:

25 an automatic speech recognition system for generating phoneme data for audio data in the data file;
a word decoder for identifying possible words within the phoneme data generated by the automatic speech recognition system; and

30 generating means for generating annotation data by combining the generated phoneme data and the decoded words.

35 35. An apparatus for generating annotation data for use in annotating a data file comprising text data, the

apparatus comprising:

a text to phoneme converter for generating phoneme data for text data in the data file; and

generating means for generating annotation data by combining the phoneme data and words in the text data.

36. An apparatus for generating annotation data for use in annotating a data file, the apparatus comprising:

input means for receiving an input voice signal;

speech recognition means for converting the input voice signal into phoneme data and words; and

generating means for generating annotation data by combining the phoneme data and the words.

37. An apparatus for generating annotation data for use in annotating a data file, the apparatus comprising:

input means for receiving a typed input from a user;

converting means for converting words in the typed input into phoneme data; and

generating means for generating annotation data by combining the phoneme data and words in the typed input.

38. An apparatus for generating annotation data for use in annotating a data file, the apparatus comprising:

means for receiving image data representative of text;

character recognition means for converting said image data into text data;

converting means for converting words in the text data into phoneme data; and

generating means for generating annotation data by combining the phoneme data and words in the text data.

39. An apparatus according to any of claims 34 to 38, wherein said annotation data defines a phoneme and word

lattice and wherein said generating means comprises:

(i) means for generating data defining a plurality of nodes within the lattice and a plurality of links connecting the nodes within the lattice;

5 (ii) means for generating data associating a plurality of phonemes of the phoneme data with a respective plurality of links; and

(iii) means for generating data associating at least one of the words with at least one of said links.

10

40. An apparatus according to claim 39, wherein said generating means is operable to generate said data defining said phoneme and word lattice in blocks of said nodes.

15

41. An apparatus according to claim 39 or 40, wherein said generating means is operable to generate data defining time stamp information for each of said nodes.

20

42. An apparatus according to claim 41, wherein said generating means is arranged to generate said phoneme and word lattice data in blocks of equal time duration.

25

43. An apparatus according to claim 40, 41 or 42, wherein said generating means is operable to generate data which defines each block's location within a database.

30

44. An apparatus according to claim 41 or any claim dependent thereon, wherein said data file includes a time sequential signal, and wherein said generating means is operable to generate time stamp data which is time synchronised with said time sequential signal.

35

45. An apparatus according to claim 44, wherein said

time sequential signal is an audio and/or video signal.

5 46. An apparatus according to claim 34 or any claim dependent thereon, wherein said audio data includes audio data which defines the parol of a plurality of speakers, and wherein said generating means is operable to generate data which defines separate phoneme and word annotation data for the parol of each speaker.

10 47. An apparatus according to claim 35 or any claim dependent thereon, wherein said text data defines the parol of a plurality of speakers, and wherein said generating means is operable to generate data defining separate phoneme and word annotation data for the parol
15 of each speaker.

20 48. An apparatus according to claim 34 or any claim dependent thereon, wherein said speech recognition system is operable to generate data defining a weighting for the phonemes in the phoneme data.

25 49. An apparatus according to claim 34 or any claim dependent thereon, wherein said word decoder is operable to generate data defining a weighting for the words identified within said phoneme data.

30 50. An apparatus according to claim 39 or any claim dependent thereon, wherein said means for generating data defining a plurality of nodes and a plurality of links is operable to define at least one node which is connected to a plurality of other nodes by a plurality of links.

35 51. An apparatus according to claim 50, wherein at least one of said plurality of links connecting said node to said plurality of other nodes is associated with a

phoneme and wherein at least one of said links connecting said node to said plurality of other nodes is associated with a word.

- 5 52. An apparatus according to claim 36 or any claim dependent thereon, wherein said speech recognition means is operable to generate data defining a weighting for the phonemes in the phoneme data.
- 10 53. An apparatus according to claim 52, wherein said speech recognition means is operable to generate data defining a weighting for the words within the word data.
- 15 54. An apparatus according to claim 36 or 37 or any claim dependent thereon, further comprising means for associating said annotation data with said data file.
- 20 55. An apparatus according to claim 37 or any claim dependent thereon, wherein said converting means comprises an automatic phonetic transcription unit which generates said phoneme data from words within the typed input.
- 25 56. An apparatus according to claim 38 or any claim dependent thereon, wherein said converting means comprises an automatic phonetic transcription unit which generates said phoneme data from words within the text data output by said character recognition means.
- 30 57. An apparatus according to claim 38 or any claim dependent thereon, further comprising means for associating said annotation data with either said image data representative of said text or with said text data.
- 35 58. An apparatus according to claim 38 or any claim

dependent thereon, wherein said receiving means comprises a document scanner or a facsimile machine.

5 59. A method of generating annotation data for use in annotating a data file comprising audio data, the method comprising the steps of:

using an automatic speech recognition system to generate phoneme data for audio data in the data file;

10 using a word decoder to identify possible words within the phoneme data generated by the automatic speech recognition system; and

generating annotation data by combining the generated phoneme data and the decoded words.

15 60. A method of generating annotation data for use in annotating a data file comprising text data, the method comprising the steps of:

using a text to phoneme converter to generate phoneme data for text data in the data file; and

20 generating annotation data by combining the phoneme data and words in the text data.

25 61. A method of generating annotation data for use in annotating a data file, the method comprising the steps of:

receiving an input voice signal;

processing the input voice signal using a speech recognition system to generate phoneme data and word data for the input voice signal; and

30 generating annotation data by combining the phoneme data and the word data generated for the input voice signal.

35 62. A method of generating annotation data for use in annotating a data file, the method comprising the steps

of:

receiving a typed input;
converting words in the typed input into phoneme
data; and
5 generating annotation data by combining the phoneme
data and words in the typed input.

63. A method of generating annotation data for use in
annotating a data file, the method comprising the steps
10 of:

receiving image data representative of text;
converting said image data into text data using a
character recognition unit;
converting words in the text data into phoneme data;
15 and
generating annotation data by combining the phoneme
data and words within the text data.

64. A method according to any of claims 59 to 63,
20 wherein said annotation data defines a phoneme and word
lattice and wherein said generating step comprises the
steps of:

(i) generating data defining a plurality of nodes
within the lattice and a plurality of links connecting
25 the nodes within the lattice;

(ii) generating data associating a plurality of
phonemes of the phoneme data with a respective plurality
of links; and

(iii) generating data associating at least one of
30 the words with at least one of said links.

65. A method according to claim 64, wherein said
generating step generates said data defining said phoneme
and word lattice in blocks of said nodes.

66. A method according to claim 64 or 65, wherein said generating step generates data defining time stamp information for each of said nodes.

5 67. A method according to claim 66, wherein said generating step generates said phoneme and word lattice data in blocks of equal time duration.

10 68. A method according to claim 65, 66 or 67, wherein said generating step generates data which defines each block's location within a database.

15 69. A method according to claim 66 or any claim dependent thereon, wherein said data file includes a time sequential signal, and wherein said generating step generates time stamp data which is time synchronised with said time sequential signal.

20 70. A method according to claim 69, wherein said time sequential signal is an audio and/or video signal.

25 71. A method according to claim 59 or any claim dependent thereon, wherein said audio data includes audio data which defines the parol of a plurality of speakers, and wherein said generating step generates data which defines separate phoneme and word annotation data for the parol of each speaker.

30 72. A method according to claim 60 or any claim dependent thereon, wherein said text data defines the parol of a plurality of speakers, and wherein said generating step generates data defining separate phoneme and word annotation data for the parol of each speaker.

35 73. A method according to claim 59 or any claim

dependent thereon, wherein said speech recognition system generates data defining a weighting for the phonemes associated with said links.

5 74. A method according to claim 59 or any claim dependent thereon, wherein said word decoder generates data defining a weighting for the words associated with said links.

10 75. A method according to claim 64 or any claim dependent thereon, wherein said step of defining a plurality of nodes and a plurality of links defines at least one node which is connected to a plurality of other nodes by a plurality of links.

15 76. A method according to claim 75, wherein at least one of said plurality of links connecting said node to said plurality of other nodes is associated with a phoneme and wherein at least one of said links connecting said node to said plurality of other nodes is associated with a word.

20 77. A method according to claim 61 or any claim dependent thereon, wherein said speech recognition system generates data defining a weighting for the phonemes associated with said links.

25 78. A method according to claim 61 or any claim dependent thereon, wherein said speech recognition system generates data defining a weighting for the words associated with said links.

30 79. A method according to claim 61 or 62 or any claim dependent thereon, further comprising the step of associating said annotation data with said data file.

35

80. A method according to claim 62 or any claim dependent thereon, wherein said converting step uses an automatic phonetic transcription unit which generates said phoneme data for words within the typed input.

5

81. A method according to claim 63 or any claim dependent thereon, wherein said step of converting words into phonemes uses an automatic phonetic transcription unit which generates said phoneme data for words within the text data output by said character recognition unit.

10

82. A method according to claim 63 or any claim dependent thereon, further comprising the step of associating said annotation data with either said received image data or with said text data.

15

83. A method according to claim 63 or any claim dependent thereon, wherein said receiving step uses a document scanner or a facsimile machine.

20

84. A method of searching a data file including annotation data in response to an input query, the method comprising the steps of:

generating phoneme data and word data corresponding to the input query;

25

searching the data file based on the phoneme data and/or the word data and the annotation data; and

outputting search results in dependence upon the result of said searching step.

30

85. A method according to claim 84, wherein said annotation data defines a phoneme and word lattice comprising:

(i) data for defining a plurality of nodes within the lattice and a plurality of links connecting the nodes

35

within the lattice;

(ii) data for associating a plurality of phonemes of the phoneme data with a respective plurality of links; and

5 (iii) data for associating at least one word with at least one of said links.

86. A method for storing a data file into a database, the method comprising the steps of:

10 combining the data file with annotation data corresponding to the data file, the annotation data including phoneme data; and

storing the data file with the annotation data.

15 87. An apparatus for searching a data file including annotation data, in response to an input query, the apparatus comprising:

means for generating phoneme data and word data corresponding to the input query;

20 means for searching a data file based on the phoneme data and/or the word data and the annotation data; and

means for outputting a search result in dependence upon the result of said searching means.

25 88. An apparatus according to claim 87, wherein said annotation data defines a phoneme and word lattice, and comprises:

(i) data defining a plurality of nodes within the lattice and a plurality of links connecting the nodes within the lattice;

30 (ii) data associating a plurality of phonemes of the phoneme data with a respective plurality of links; and

(iii) data associating at least one word with at least one of said links.

89. An apparatus for storing a data file into a database, the apparatus comprising:

means for inputting the data file and annotation data corresponding to the data file, the annotation data including phoneme data; and

means for storing the data file with the annotation data.

90. A medium for storing a data file, the data file comprising:

an audio data; and

an annotation data corresponding to the audio data, said annotation data including phoneme data.

91. A medium for storing a data file, the data file comprising:

video data;

audio data corresponding to the video data; and

annotation data corresponding to the audio data, the annotation data including phoneme data.

92. A medium for storing a data file, the data file comprising:

text data; and

annotation data corresponding to the text data, said annotation data including phoneme data.


93. Data including audio data and further comprising annotation data corresponding to the audio data, which annotation data includes phoneme data.

94. Data including video data and further comprising audio data corresponding to the video data and annotation data corresponding to the audio data, which annotation data includes phoneme data.

95. Data including text data, the data further comprising annotation data corresponding to the text data, which annotation data includes phoneme data.

5 96. A data carrier carrying data according to any of claims 1 to 13 or processor implementable instructions for controlling a processor to implement the method of any one of claims 14 to 23 or 59 to 83 or 84 to 86.

10 97. Processor implementable instructions for controlling a processor to implement the method of any one of claims 14 to 23 or 59 to 83 or 84 to 86.



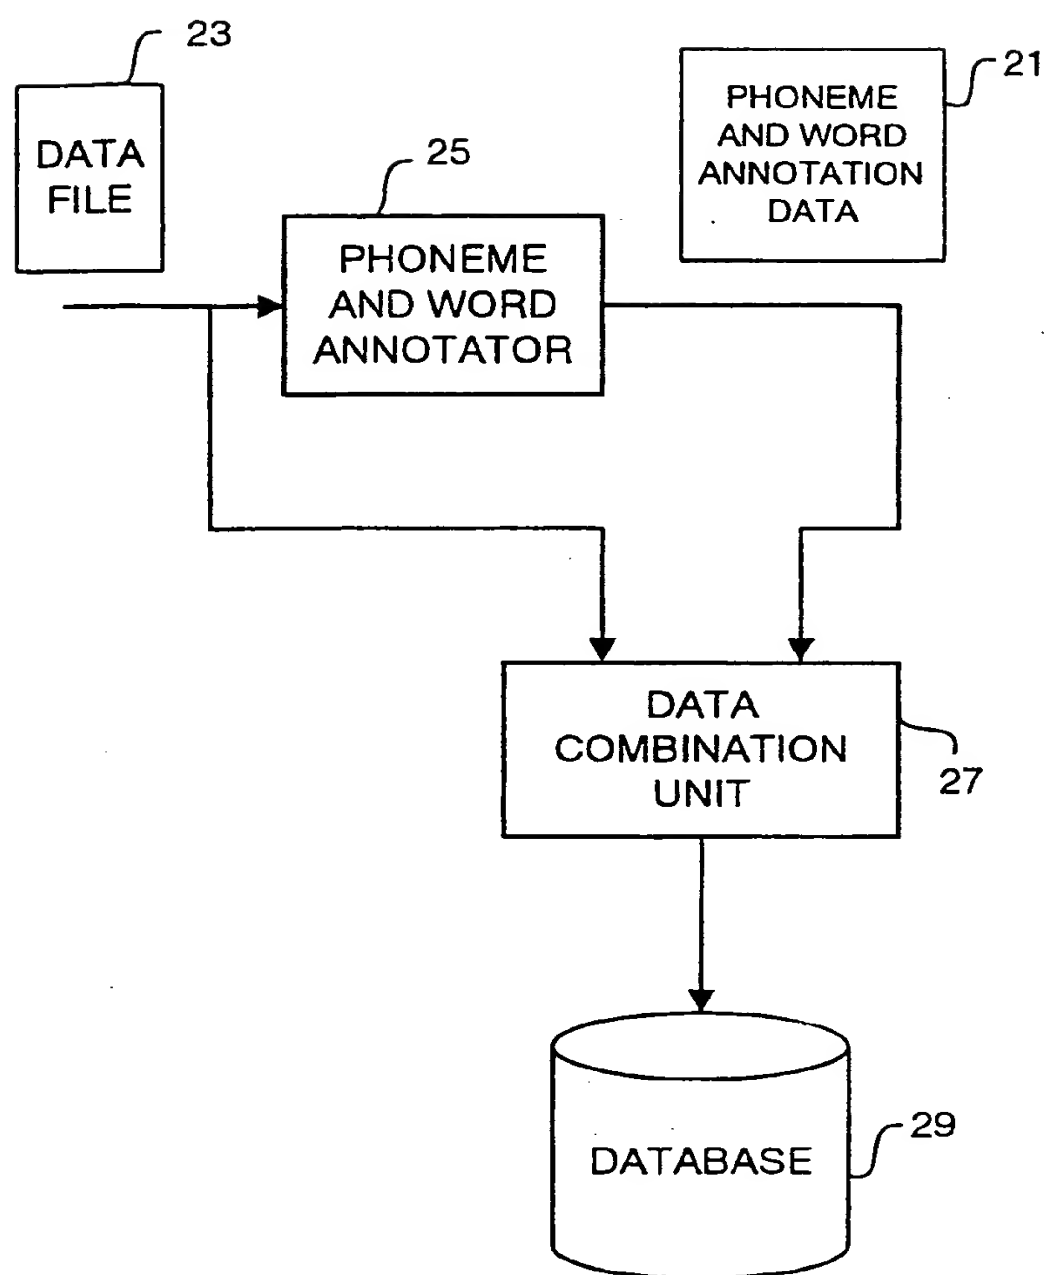


Fig. 2




Fig. 3




Fig. 4a




Fig. 4b

5/21


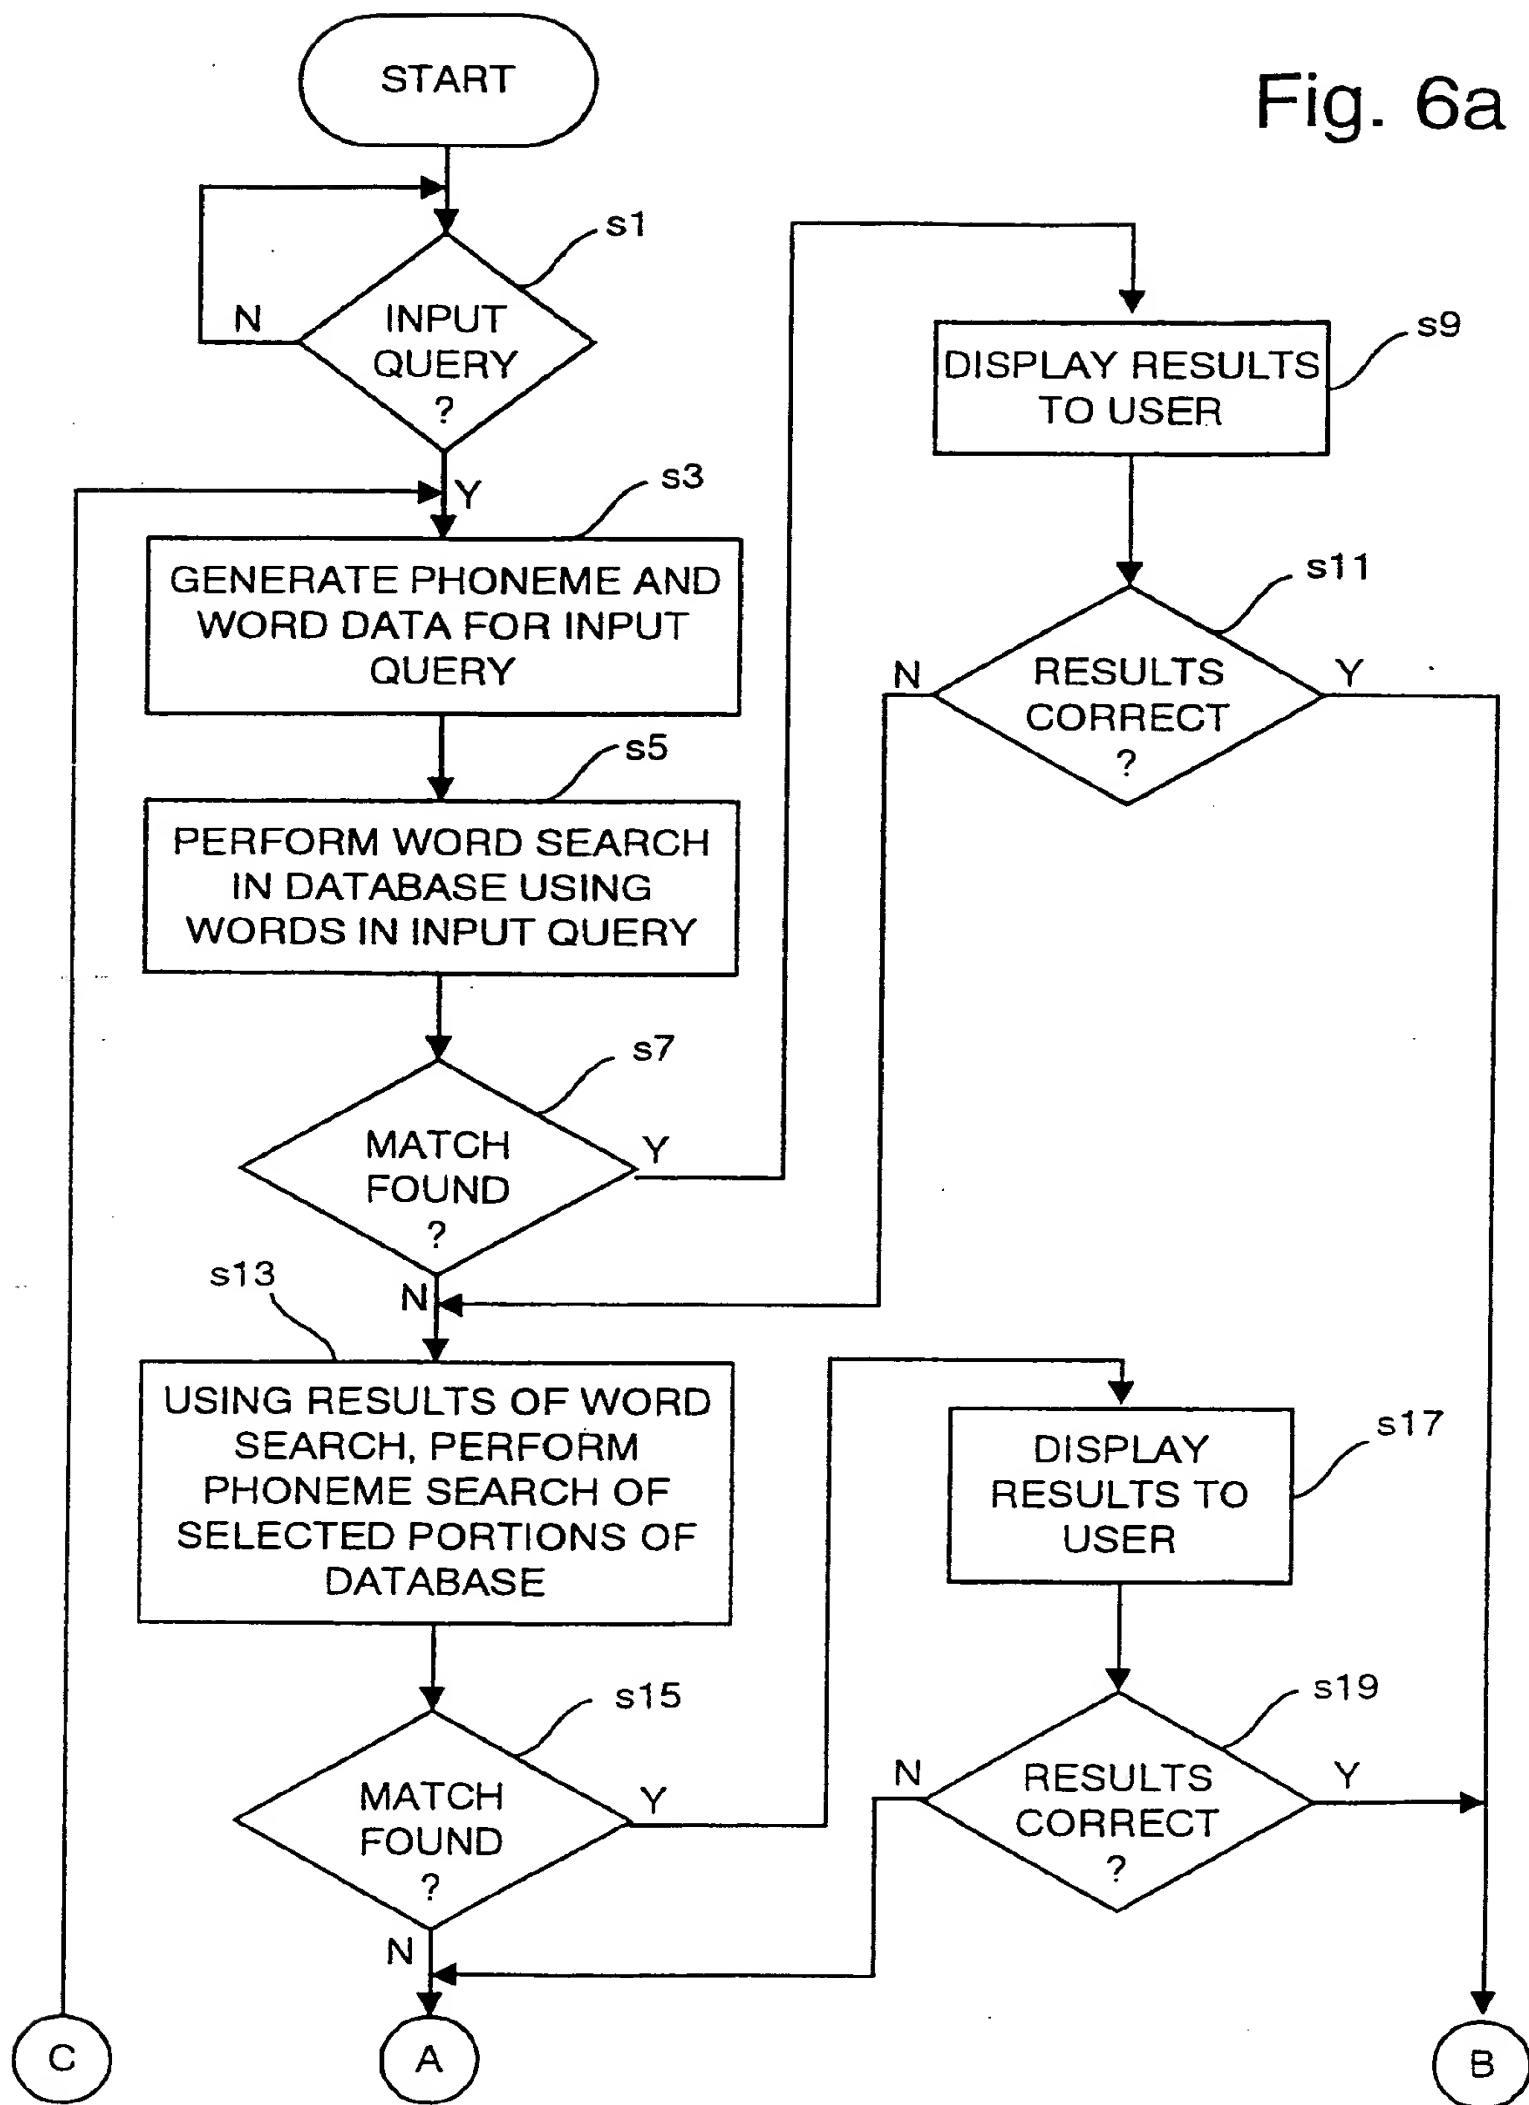


Fig. 5

6/21

Fig. 6a



7/21

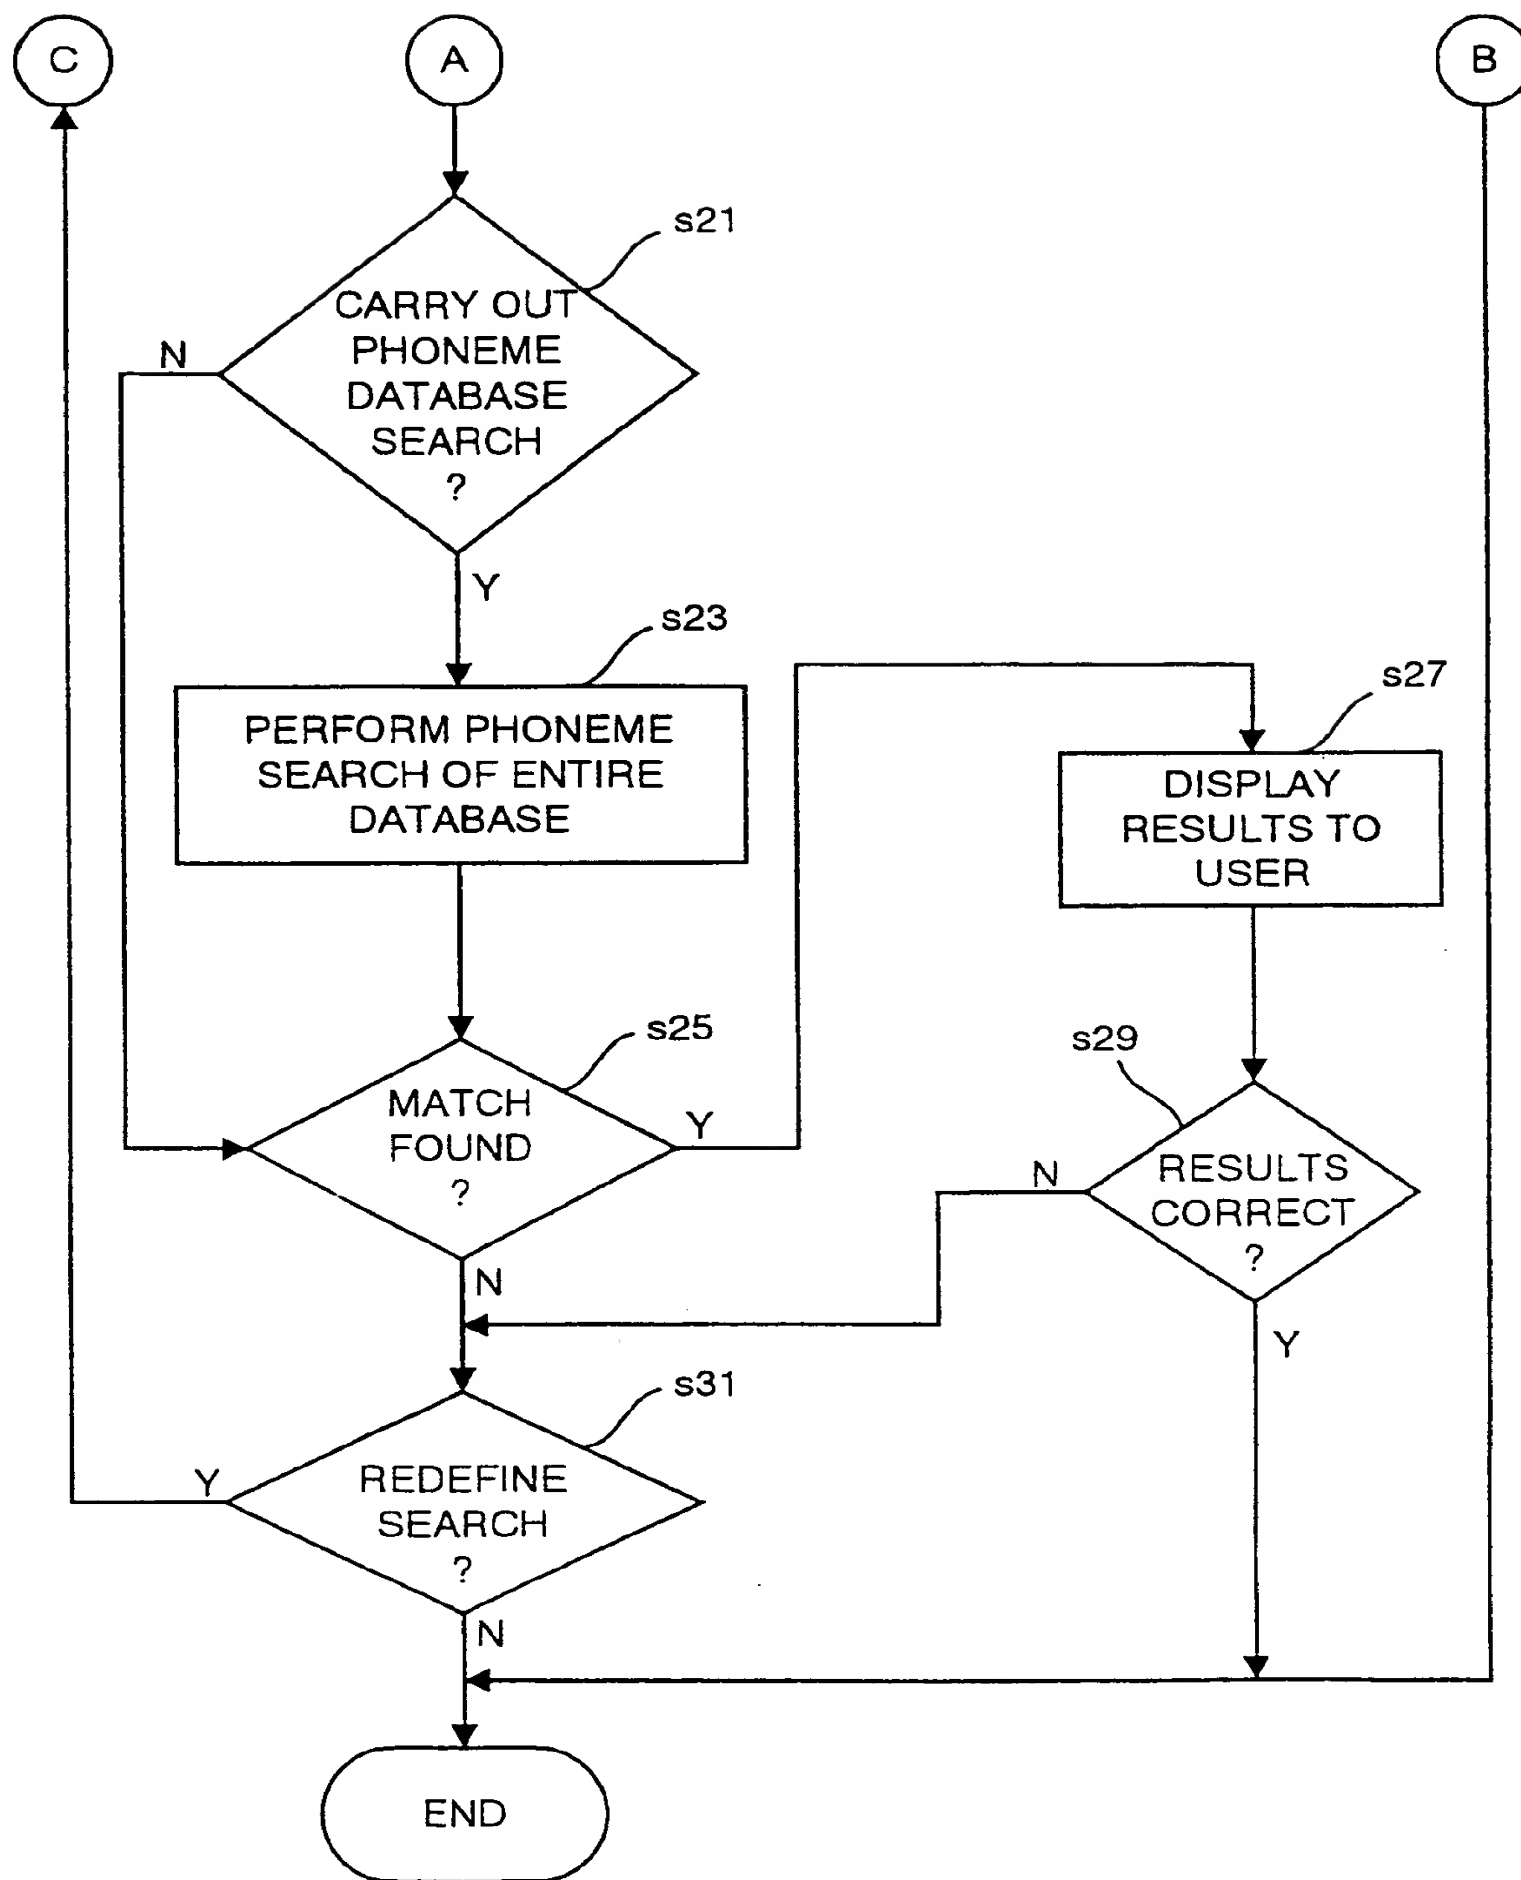


Fig. 6b

8/21

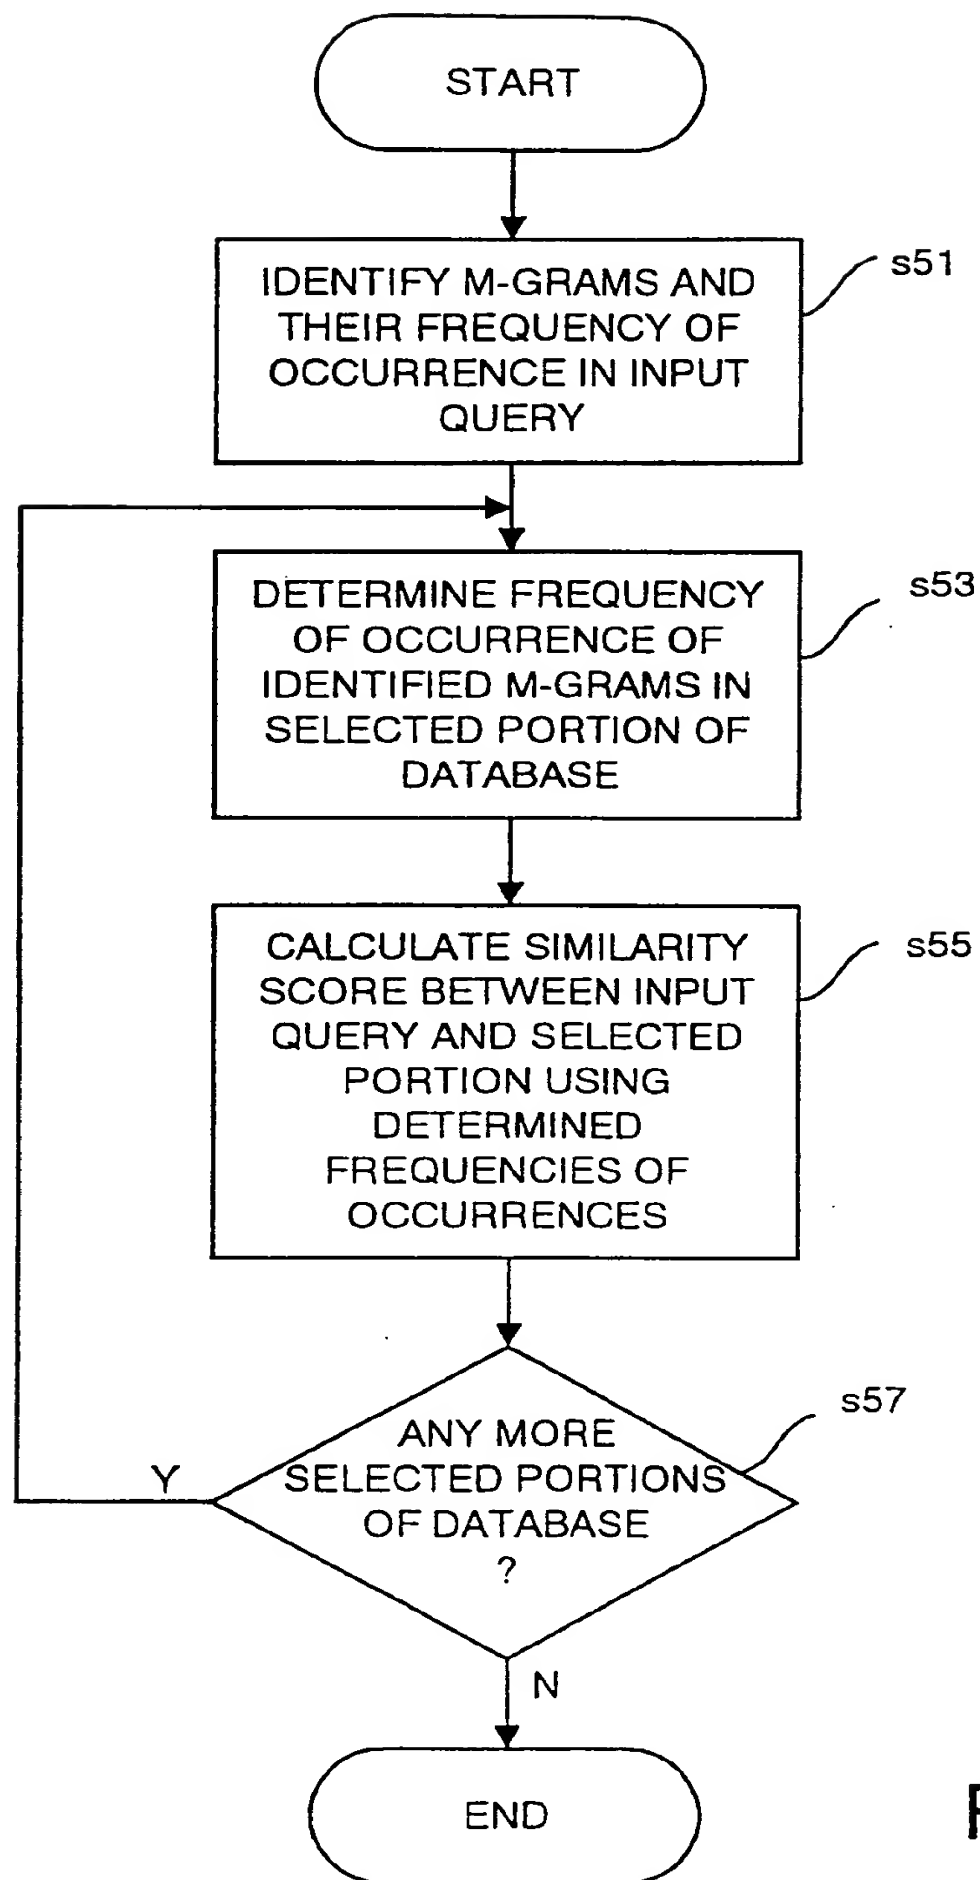


Fig. 7

9/21




Fig. 8




Fig. 9




Fig. 10

11/21




Fig. 11

12/21




Fig. 12

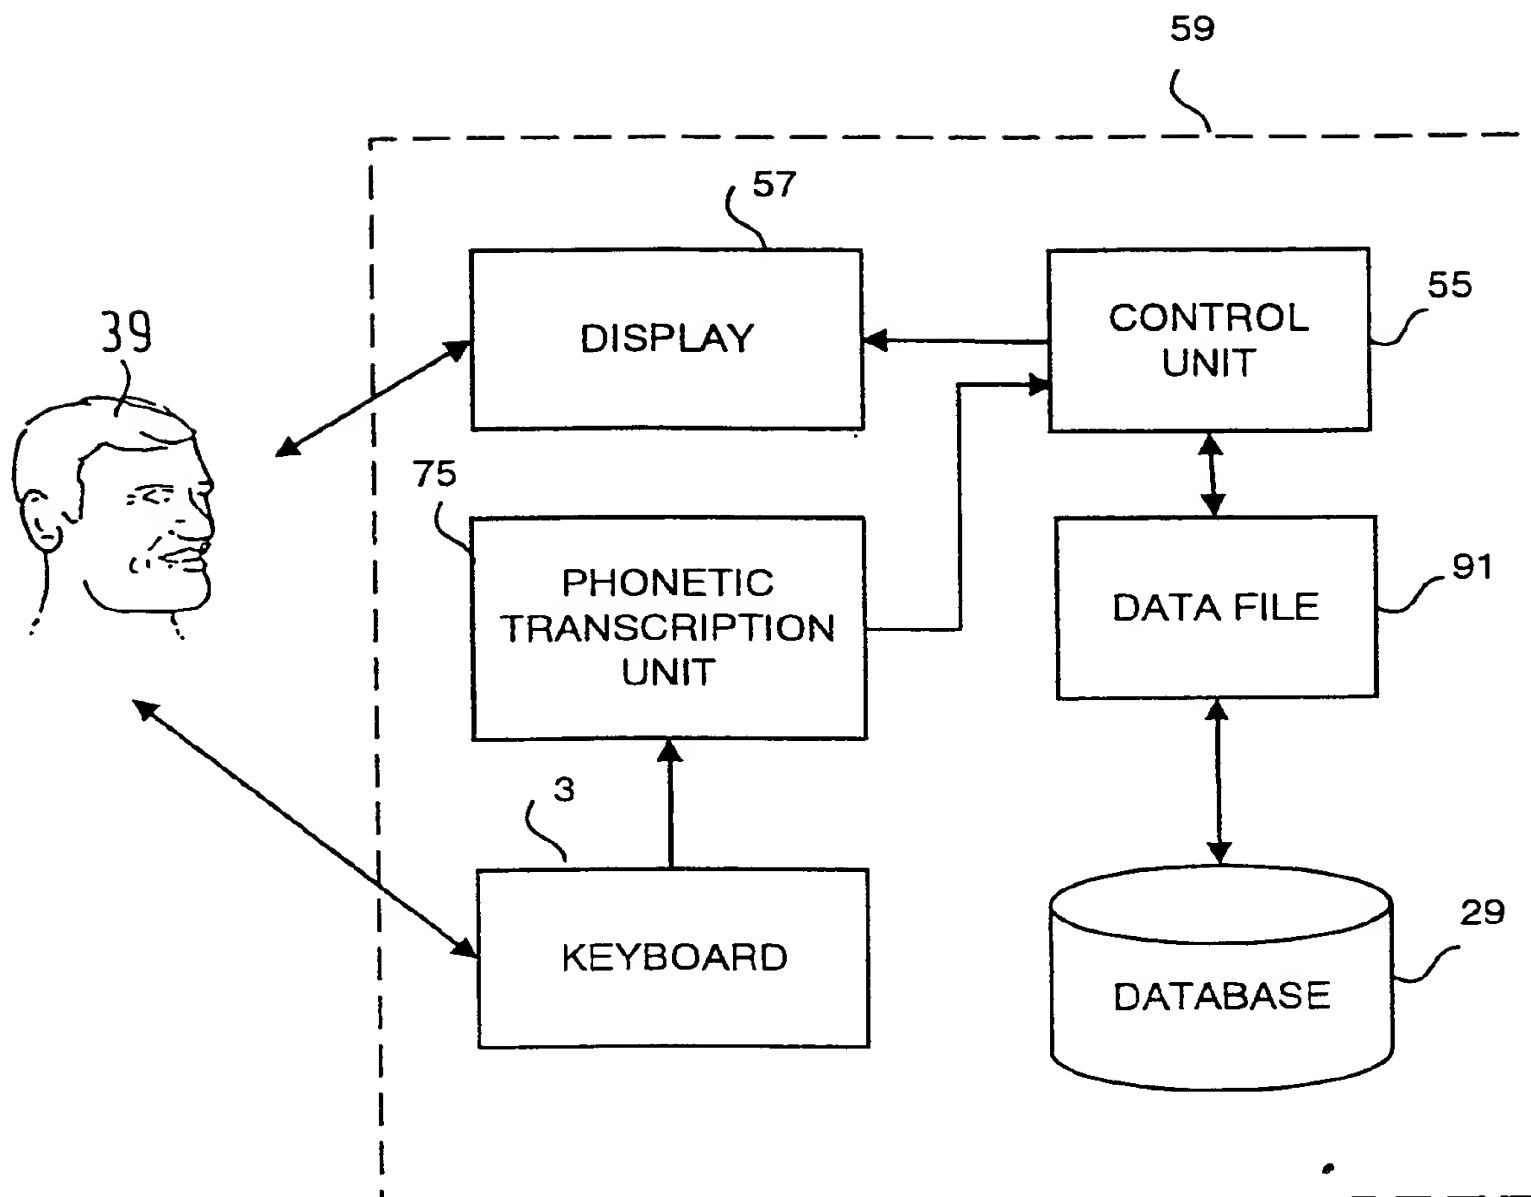



Fig. 13



14/21

Fig. 14

15/21

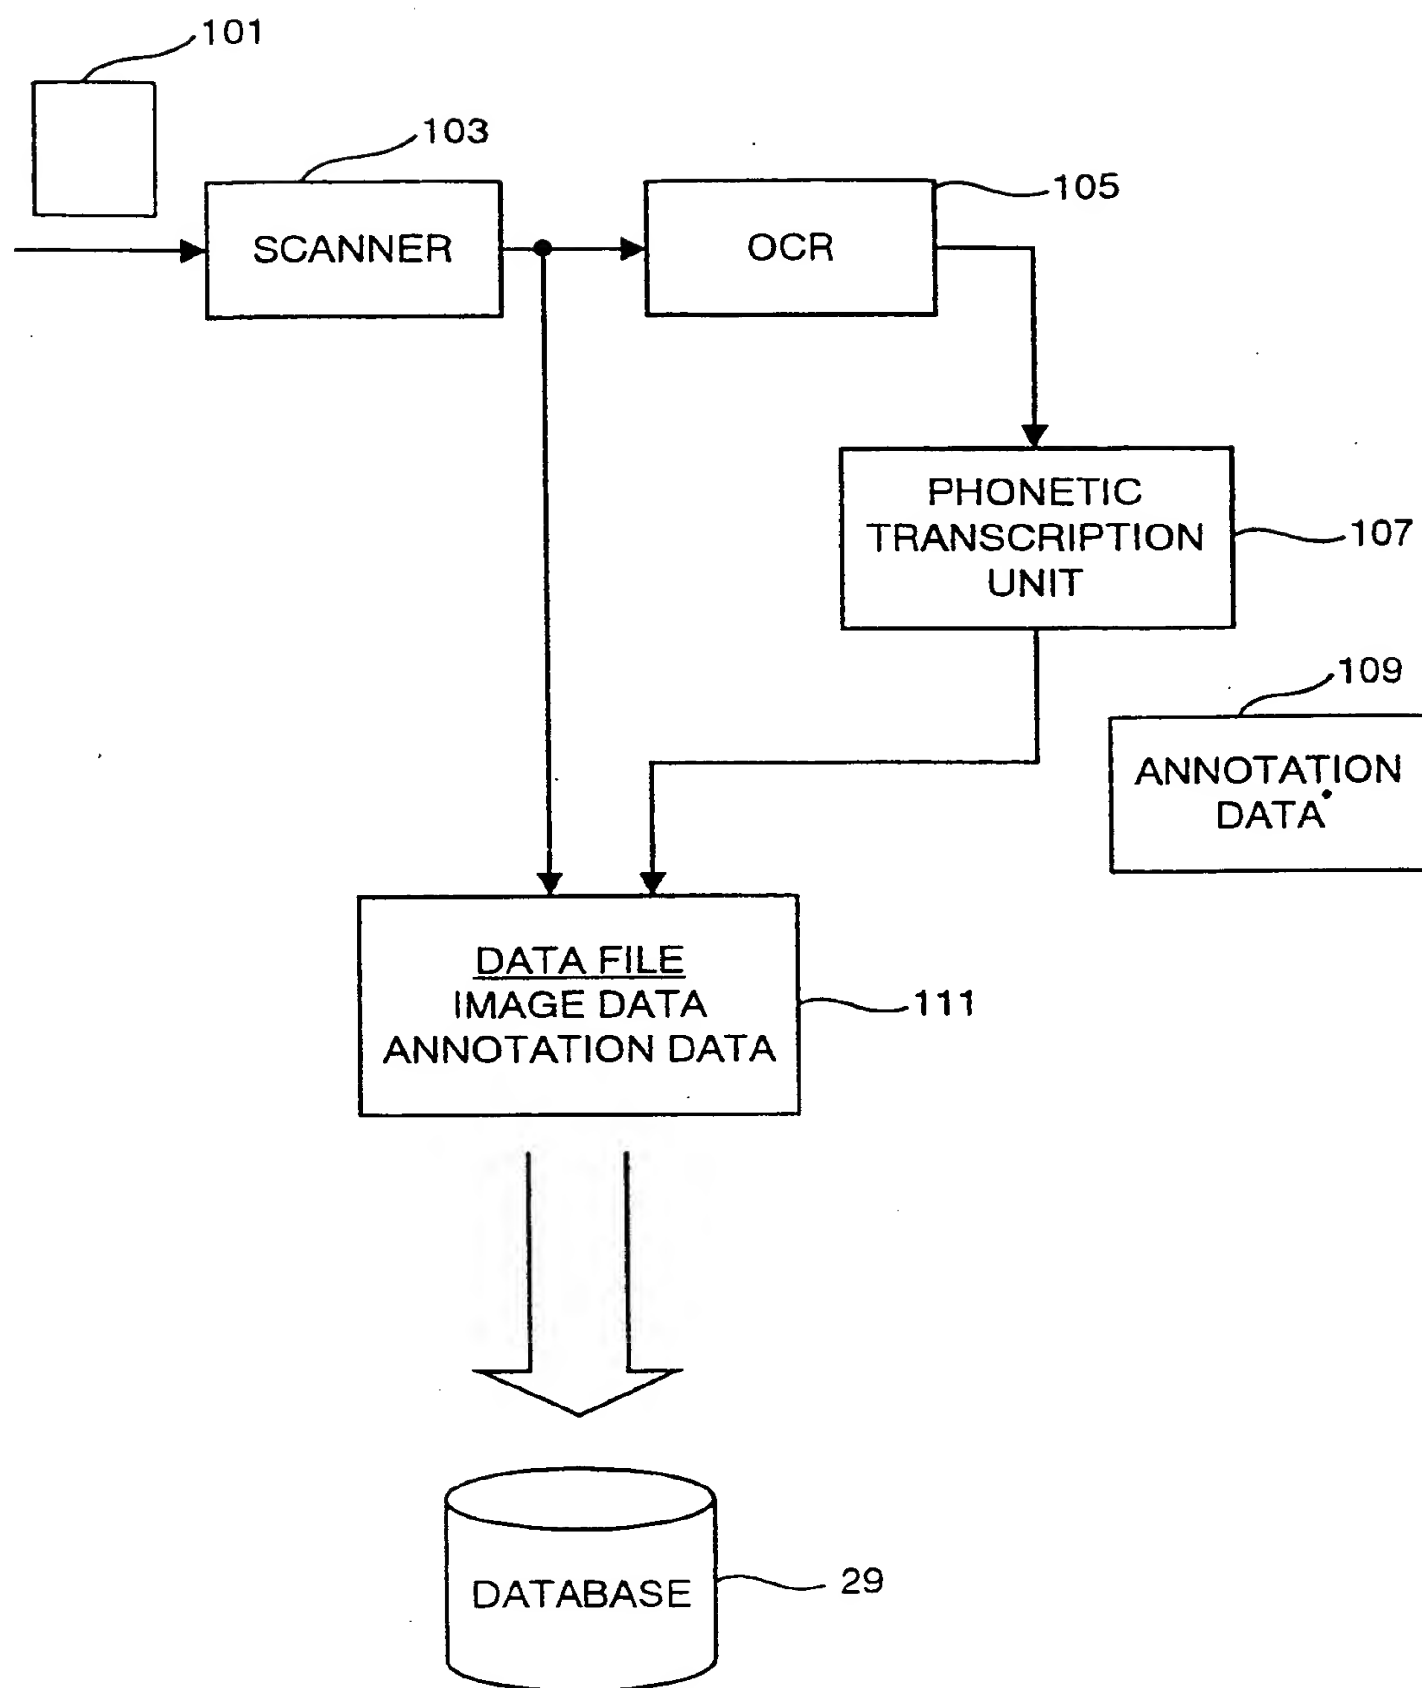


Fig. 15

16/21

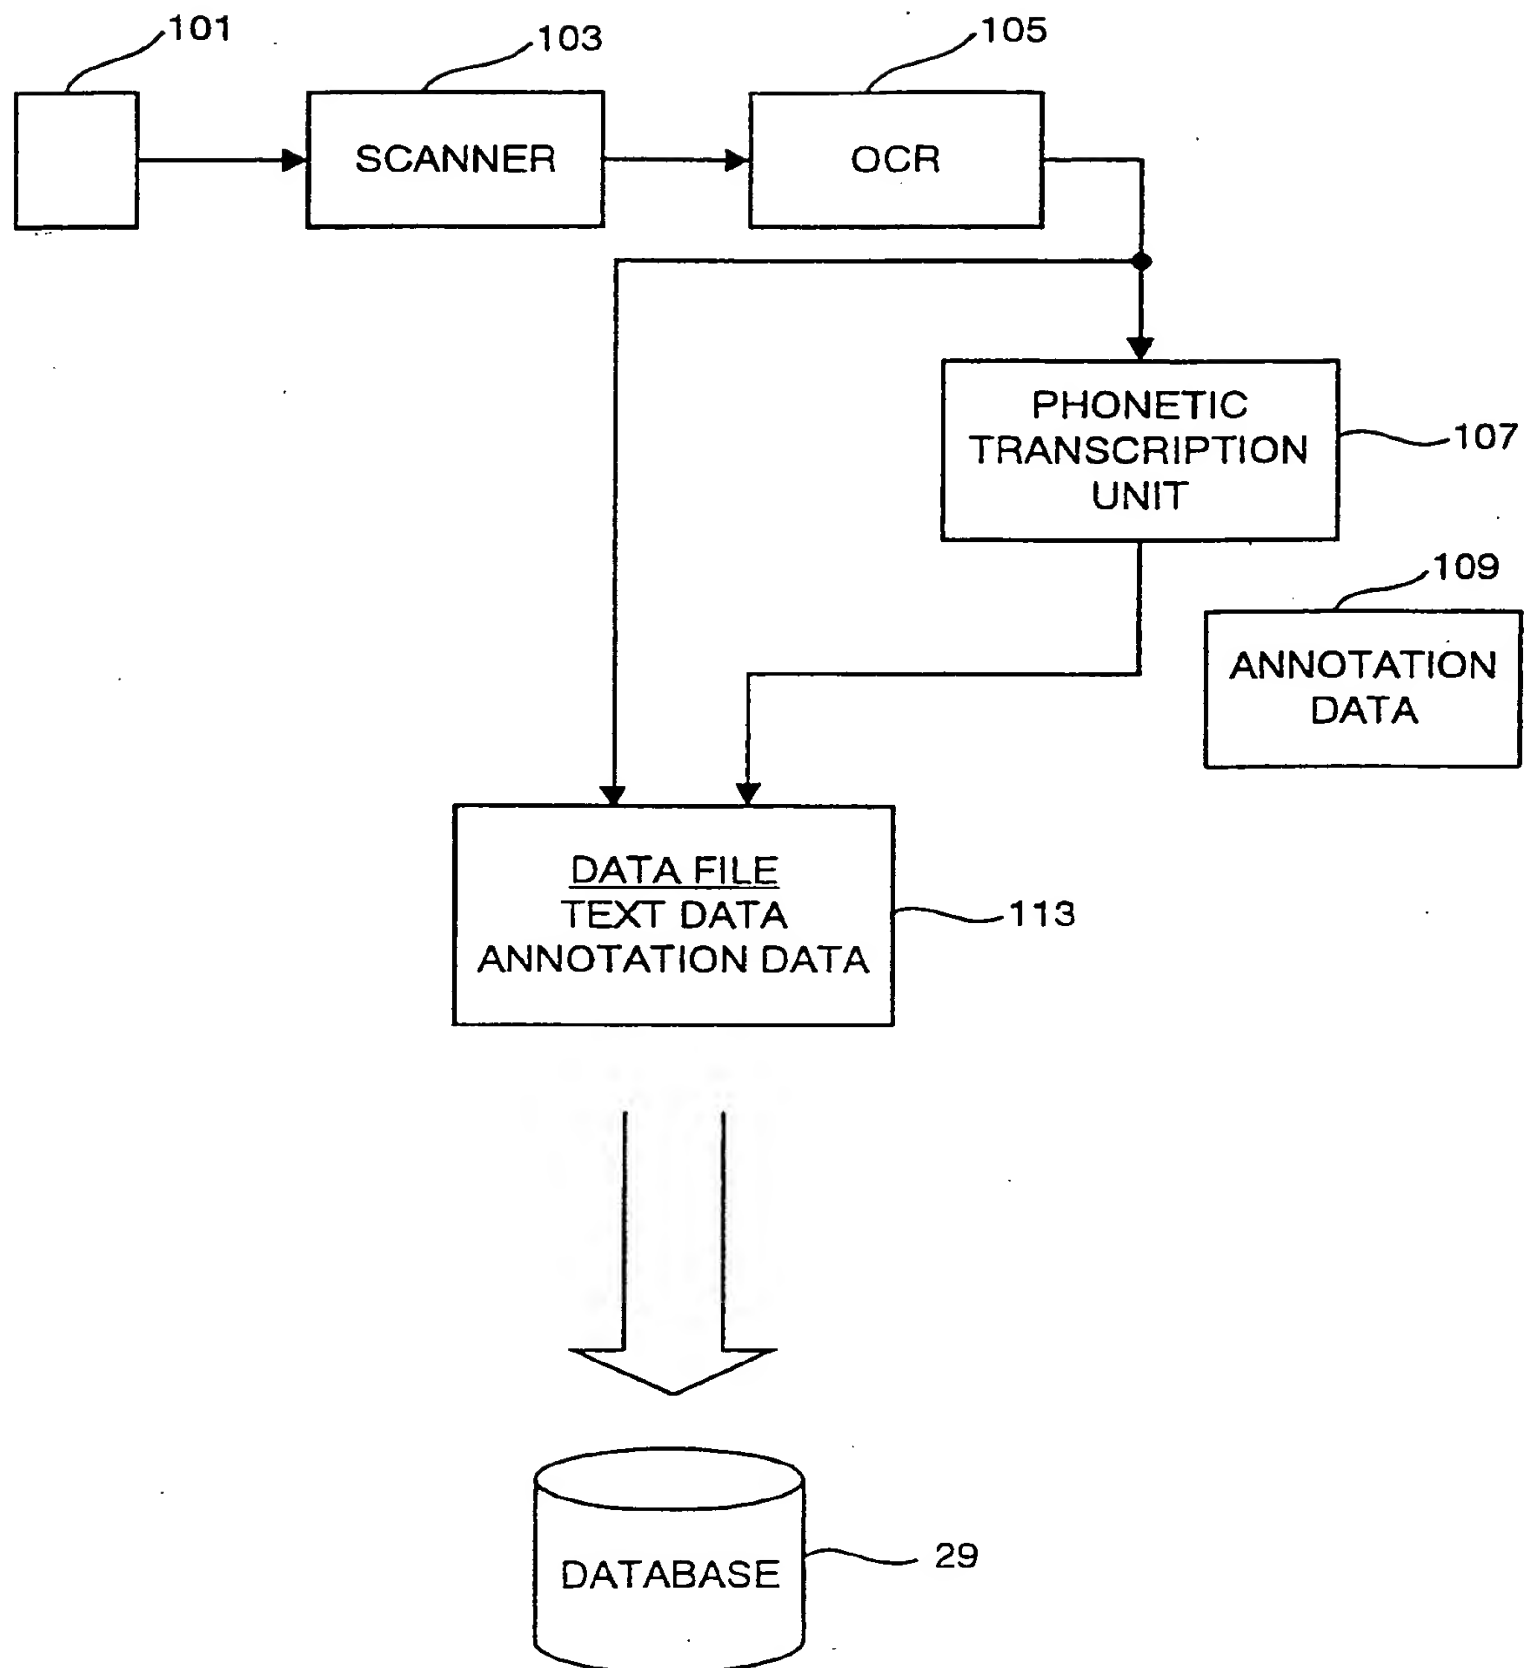


Fig. 16

17/21

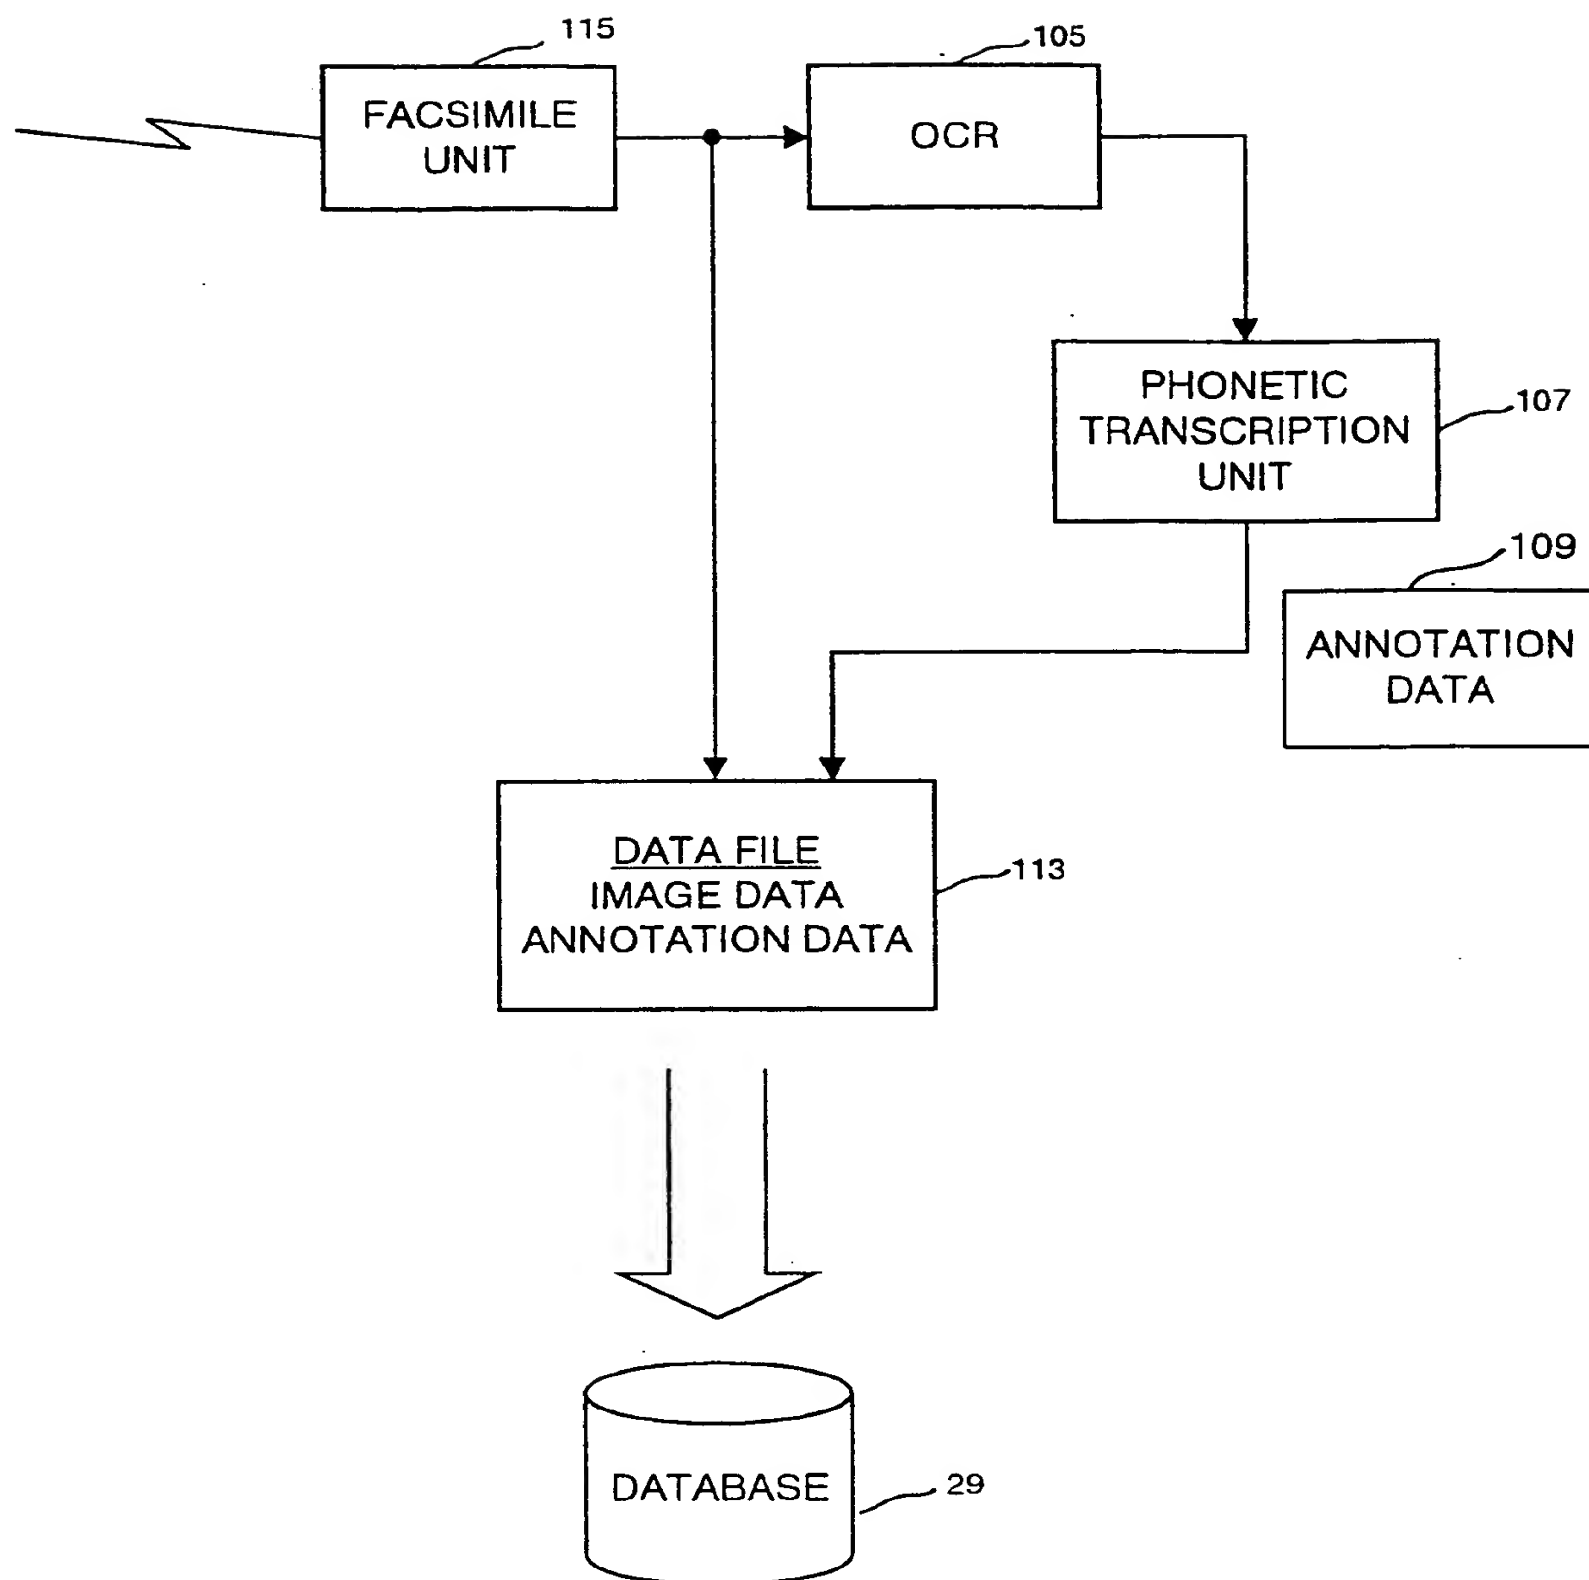


Fig. 17

18/21

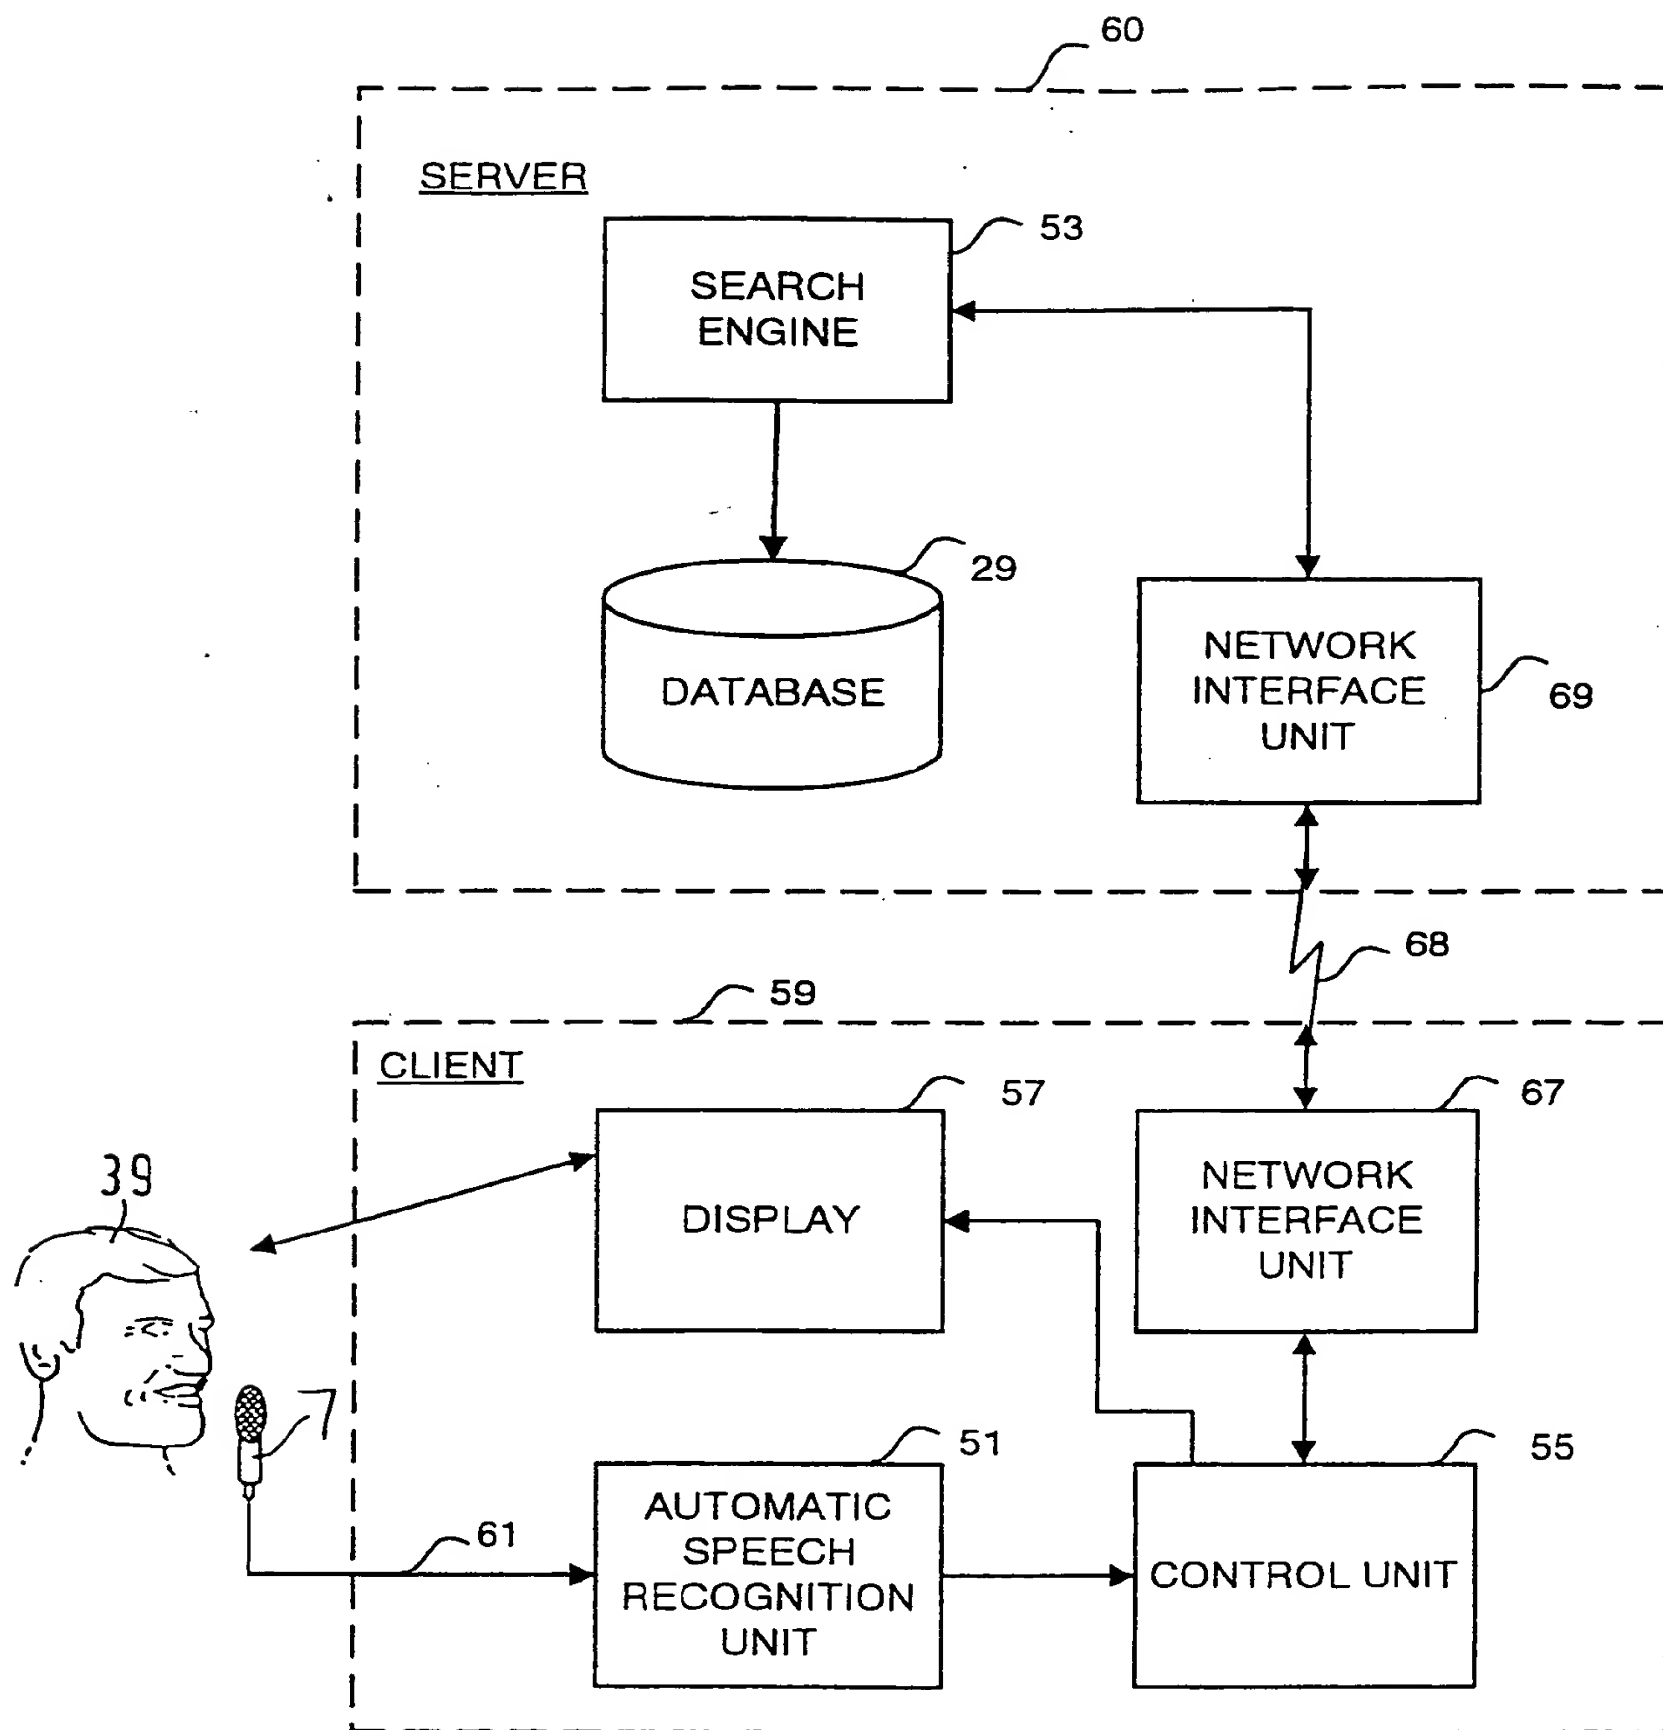


Fig. 18

19/21




Fig. 19




Fig. 20

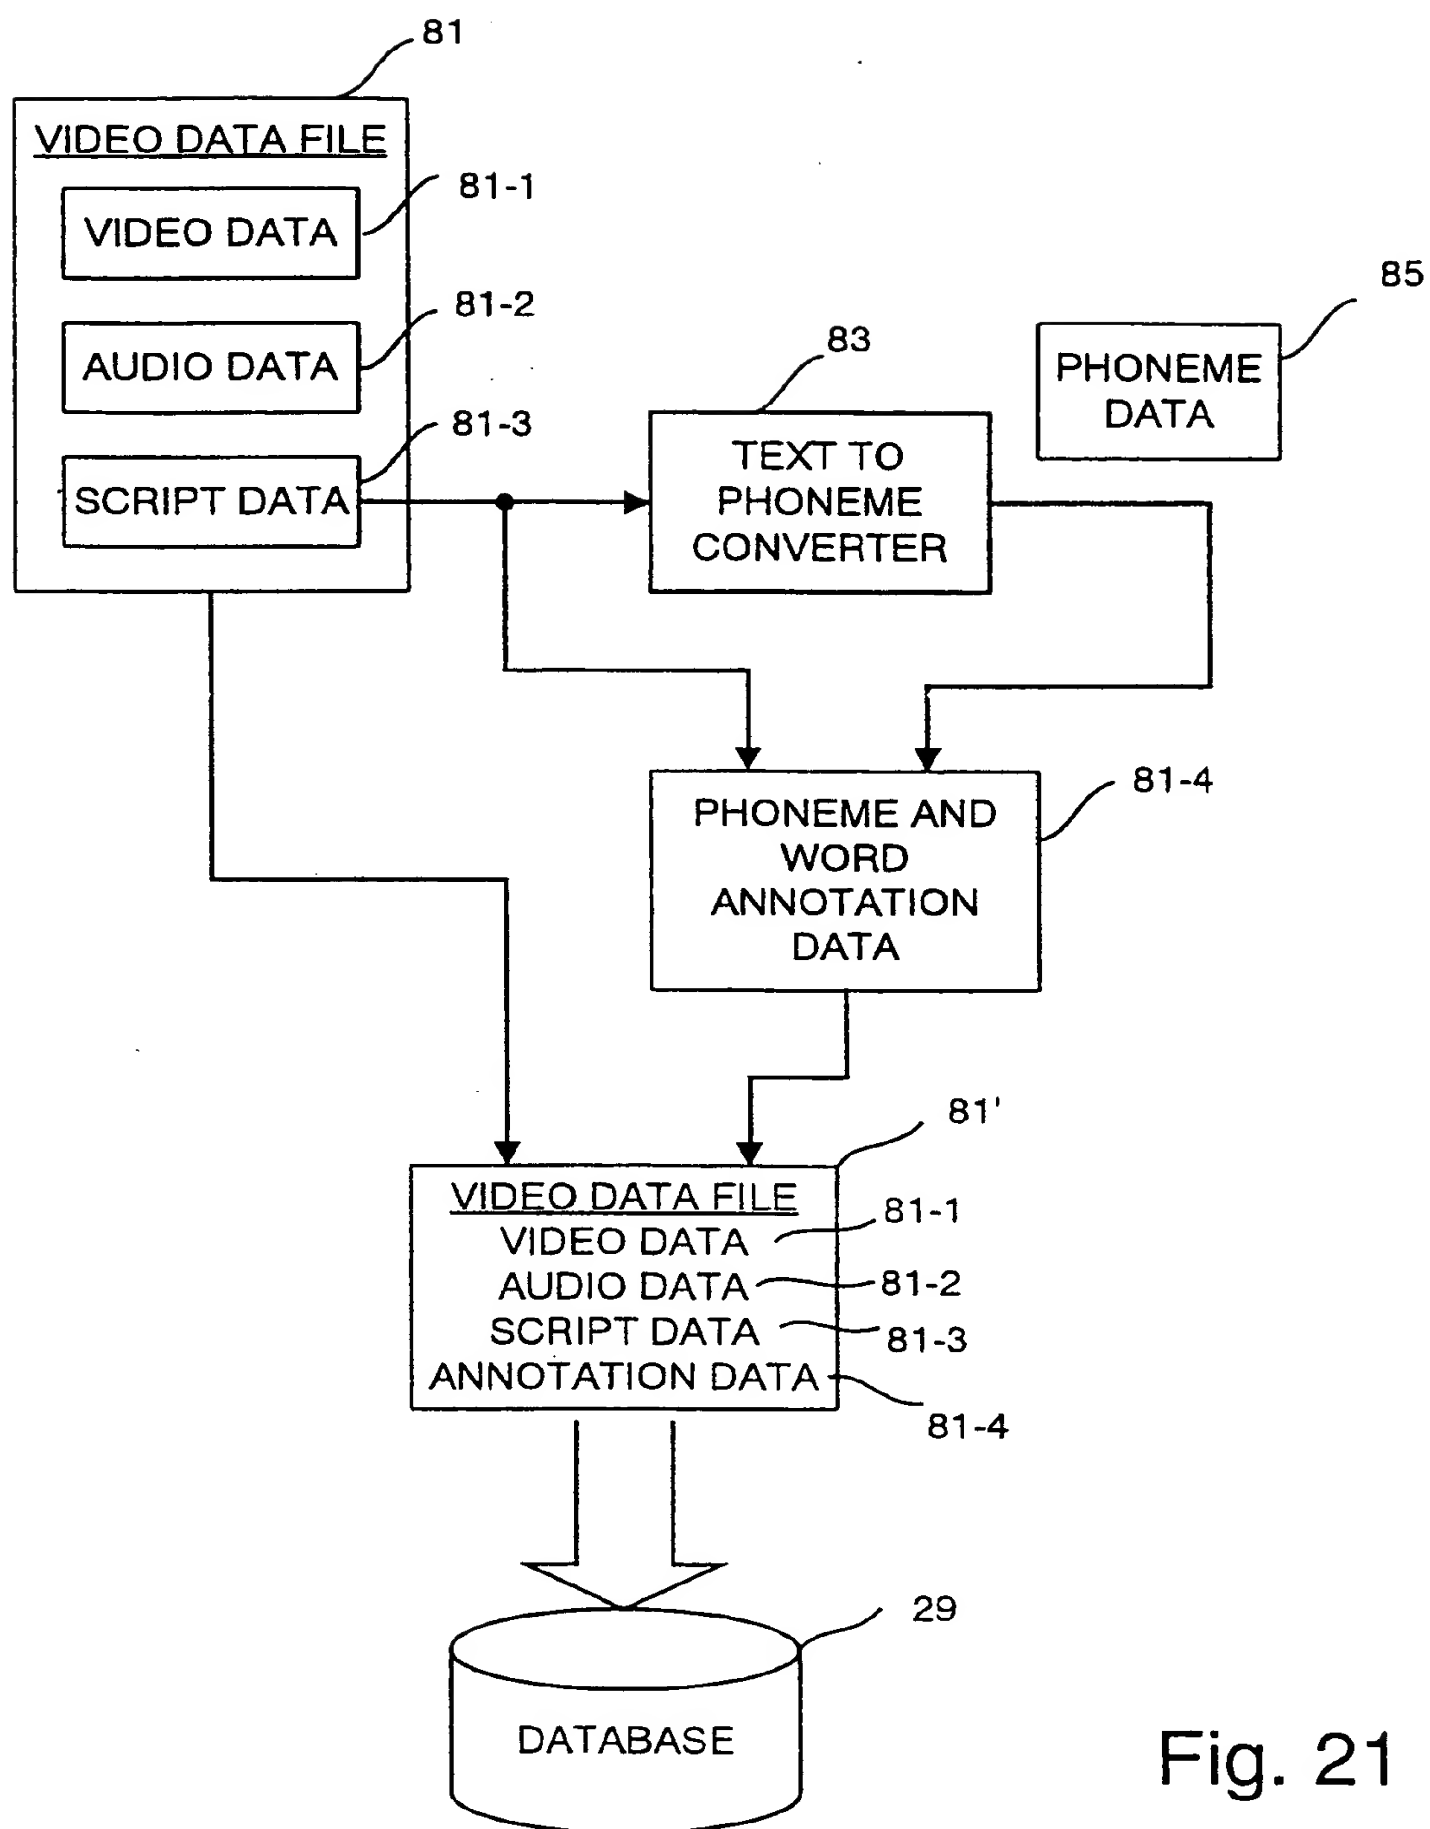


Fig. 21